Long-Run Intergenerational Effects of Social Security *

Daniel K. Fetter
Stanford University
and NBER

Lee M. Lockwood
University of Virginia
and NBER

Paul Mohnen
Univ. of Pennsylvania

November 22, 2021

(Link to latest version)

Abstract

Both historically and today, much of the support of the elderly by their adult children has taken the form of in-kind transfers that require shared location, such as shared housing. To the extent that Social Security substitutes for these forms of support, it might thereby relax a "location constraint," with potentially wide-ranging intergenerational effects. We investigate this issue by combining a novel empirical approach—exploiting within-occupation, cross-industry differences in Social Security coverage in the early years of the program—with a dataset linking information on parents to the long-run outcomes of their children. We find that individuals whose parents were more likely to have received Social Security, or received it earlier, were less likely to live in their parents' location, earned more, and lived in ZIP codes of higher socioeconomic status near the end of their lives. A variety of evidence suggests that migration to better-matched labor markets drove most of these gains. The magnitude of the estimates suggests that the impact of government old-age support programs on the allocation of people across locations may be central to their overall welfare effects.

*For helpful comments and discussions, we thank Ran Abramitzky, Martha Bailey, Amy Finkelstein, Matt Notowidigdo, Heidi Williams, as well as seminar participants at Harvard, the University of Michigan, Stanford University, and the NBER Summer Institute. We thank Tamri Matiashvili, Elisa Heinrich Mora, Maxwell Rong, Amanda Wahlers, and Qiyi Zhao for excellent research assistance. Mohnen gratefully acknowledges financial support from the Alfred P. Sloan Foundation Pre-Doctoral Fellowship on the Economics of an Aging Workforce, awarded through the NBER. Fetter and Lockwood gratefully acknowledge financial support from National Science Foundation (NSF) Grant SES-1628860. Fetter gratefully acknowledges financial support from the Stanford Institute for Economic Policy Research, and Lockwood gratefully acknowledges financial support from the Bankard Fund for Political Economy.

1 Introduction

There is considerable evidence that government old-age support programs, like Social Security, affect intergenerational transfers between adult children and their elderly parents (e.g., Costa, 1997, 1999; McGarry and Schoeni, 2000; Engelhardt et al., 2005). That such changes in family transfers can fundamentally transform the effects of such programs is well known (e.g., Barro, 1974; Becker, 1974; Bernheim and Bagwell, 1988). Yet leading models, which have guided subsequent empirical work, focus on a special case in which family support is, at least implicitly, in the form of frictionless cash transfers. Both historically and today, however, a large share of family support has been in kind, such as shared housing, assistance with household chores, and personal care. While such support has important advantages, a potential cost is that it often requires parents and their adult children to live close to one another. To the extent that Social Security reduces the need for such transfers, it might help children of recipients by effectively relaxing a constraint that they live near their parents, which could be especially helpful if their parents' location is a poor match for their skills.

In this paper, we provide new evidence on the intergenerational effects of government old-age support by studying the long-run effects of Social Security on recipients' children and the role of migration as a driving force. The introduction and subsequent expansions of the Social Security program coincided with striking changes in the connections between the elderly and their adult children. Figure 1 shows that as government old-age payments—primarily through Social Security—increased over the second half of the 20th century, the rate of intergenerational coresidence, a common measure of intergenerational connections that encompasses transfers of various kinds (e.g., Costa, 1997, 1998, 1999; McGarry and Schoeni, 2000; Ruggles, 2007), declined substantially.

Naturally, this time-series pattern does not establish a causal relationship between government-provided old-age support and family connections. To that end, we introduce a novel empirical approach for estimating the causal impact of Social Security. This approach is based on the fact that employment in some industries was covered when the program was introduced in 1935, while similar employment in other industries was covered only in 1950 or later as the program was expanded. For example, work as a janitor in a private, for-profit firm was covered in 1935, while work as a janitor in a nonprofit hospital or educational institution was not covered until 1950. As a result, due to their industry of employment, individuals with similar jobs sometimes differed in the likelihood of ever becoming eligible for Social Security, or became eligible for Social Security benefits at different ages. We leverage this variation to compare the long-run outcomes of otherwise-similar individuals whose parents had different levels of coverage from Social Security.

Implementing this approach requires having sufficient information on parents to predict their

likely eligibility for Social Security, as well as observing their children's outcomes in adulthood. In practice, datasets typically contain information on one generation but not the other. To overcome this challenge, we build a new dataset linking Social Security death records to the 1930 US Census. The 1930 Census allows us to predict the age at which parents would likely become eligible for Social Security, based on their year of birth and their employment information prior to the passage of the Social Security Act. Social Security death records contain information on the ZIP code in which their children last lived, which serves as a proxy measure for their end-of-life socioeconomic status (SES) and also allows us to study their long-term migration patterns. We supplement this dataset with the 1973 Survey of Occupational Changes in a Generation (OCG), which contains sufficient retrospective information on respondents' fathers to determine their likely Social Security eligibility age, and allows us to observe direct measures of respondents' income.

Conceptually, we emphasize that an important omission from standard models of the intergenerational consequences of government old-age support is that, in contrast to Social Security, family support is often non-monetary. Figure 1 shows that co-residence was historically quite common, and both historical and modern data show that various types of in-kind care, such as performing errands, preparing food, or providing personal care, are also common (e.g., Sussman, 1965; Horowitz, 1985; Folbre, ed, 2012). An important characteristic of these kinds of connections is that they require spatial proximity. Hence, to the extent that Social Security displaces these connections, it may affect where people live, and hence their labor market opportunities. Also important is that both contemporary and historical evidence reveals differences between sons and daughters in the amount and type of care provided (e.g., Horowitz, 1985; Folbre, ed, 2012; Grigoryeva, 2017). Daughters more often provide types of care that, especially historically, have been more difficult to purchase on the market, or that families may prefer not to purchase on the market. Hence, it is likely that sons and daughters would be affected differently by their parents' receipt of monetary transfers.

Our main empirical finding is that children whose parents were more likely to ever be eligible for Social Security benefits, or were likely eligible at a younger age, lived in ZIP codes with higher incomes, higher house prices, and better amenities, in part through greater rates of long-term migration. These effects are concentrated among children from smaller families, consistent with the idea that displacement of family support should have smaller effects on the average child from a larger family. Moreover, consistent with daughters' greater tendency to provide types of care that are not typically purchased on the market, the effects are largely driven by sons rather than daughters. Direct evidence on income from the OCG supports the interpretation that residence in higher-income ZIP codes reflects higher lifetime income for sons.

Using a variety of approaches, we quantify the financial gains to recipients' children from their parents' Social Security coverage and the role of migration as an underlying driver. We

find that Social Security conferred large, long-lasting financial gains on recipients' children. Our preferred estimate is that a ten-year difference in fathers' predicted earliest age of Social Security eligibility—roughly the difference between 1935- and 1950-covered fathers from the 1877 birth cohort—conferred a financial gain on recipients' children of $640 per year or $12,800 in expected present value (amounts in 2000 dollars). Given the long-lasting nature of these gains and that the average recipient had 4.6 children, such gains imply that the total financial gain to recipients' children likely exceeded the gain to recipients themselves.

A variety of evidence suggests that the key driver of the financial gains to recipients' children was that Social Security led some to migrate to better-matched labor markets. A ten-year difference in fathers' predicted earliest age of Social Security eligibility increased migration to a different commuting zone (CZ), a widely-used measure of local labor markets, by roughly 2 percentage points. Social Security may have also led to earlier or farther migration conditional on migrating. Plausible gains from such effects on migration, based on typical estimates in the literature, could account for our estimates of the significant gains from Social Security to recipients' children. In contrast, even aggressive estimates of the gains from *non*-migration factors would not explain even half of the gains to recipients' children. Our results suggest that gains from migration to better-matched labor markets were not only the key driver of the large financial gains from Social Security to recipients' children, they were also likely a key component of the overall efficiency effects of Social Security. Why did Social Security increase migration? Various pieces of evidence—such as the evidence of gains for sons but not daughters—are suggestive that the underlying mechanism was a displacement of transfers from children to parents that required spatial proximity.

Our results suggest that the effects of Social Security on recipients' children may be central to its overall distributional and welfare effects. While much research has documented Social Security's large statutory redistribution from younger to older cohorts (see, e.g., Moffitt, 1984; Leimer, 1994; Feldstein and Liebman, 2002), our results suggest that the economic incidence of Social Security was much less concentrated on older cohorts. Indeed, our results suggest that even Social Security's large net transfers to early recipients may have conferred larger financial gains on recipients' children than on recipients themselves. While much of the literature on Social Security and related programs has focused on the effects on recipients' labor supply and savings (e.g., Coile, 2015; Attanasio and Brugiavini, 2003), our results suggest that the effects on recipients' children are another important determinant of the overall welfare effects.

As noted earlier, the theoretical literature on Social Security has recognized the importance of potential changes in family relationships for the effects of Social Security (Barro, 1974; Becker, 1974; Bernheim and Bagwell, 1988), though the focus has been, at least implicitly, on the case of frictionless monetary transfers. In parallel, an empirical literature has emphasized that Social Security and other old-age transfer programs displaced in-kind transfers within families in the form

of shared housing (Costa, 1997, 1999; McGarry and Schoeni, 2000; Engelhardt et al., 2005), and, in more recent periods, that Social Security increased downstream monetary transfers and reduced upstream non-monetary transfers (Mukherjee, forthcoming). However, due to data limitations, little is known about how these changes have affected recipients' children in the long run, which is a key input into the overall welfare effects of Social Security. Thanks to newly available data and our ability to follow the same individuals over long periods of time, our work builds on these findings to examine long-term effects on the location of recipients' children, as well as their labor market outcomes.

Taken together, the results suggest that to an important degree, the introduction and expansion of Social Security was not so much a large transfer from one generation to another, as it was a replacement of one type of old-age support with another. This view is in accordance with that of work by historians such as Gratton (1996), who alludes to the significant burdens imposed by the family-based system and notes that in public opinion surveys, Social Security was as popular with the young as it was with the old. Our finding of quantitatively significant effects on adult children, including through greater migration, is also consistent with evidence on the impact of the introduction of pensions in modern-day developing countries.[1]

2 Background

2.1 Eligibility and benefits in the early Social Security Program

The Social Security Act was passed in 1935, but it was under the 1939 Amendments to the Social Security Act that the first monthly retirement and survivors' benefits under Old Age and Survivors Insurance (henceforth, "Social Security") were paid, in 1940. Under the 1939 and subsequent amendments, a worker became eligible to receive monthly retirement benefits by working for a sufficient number of quarters in employment covered by Social Security, provided he or she earned above a threshold earnings level in that quarter.[2] For example, under the 1939 Amendments, a worker would be eligible for retirement benefits if he or she earned a sufficient amount in covered employment in half of the quarters after 1936 (or the quarter he or she turned 21) and prior to death or age 65, with a minimum of six quarters.

The requirement for a sufficient length of employment meant that relatively few of the birth cohorts who were already advanced in age in 1940 were eligible for monthly Social Security pay-

[1] Similar themes and findings have been documented in the context of South Africa (Jensen, 2003; Edmonds et al., 2005; Ardington et al., 2009; Hamoudi and Thomas, 2014), China (Chen et al., 2018; Eggleston et al., 2018; Huang and Zhang, 2021), and Ghana and Indonesia (Bau, 2021).

[2] The original act in 1935 specified a different criterion for coverage, which required both sufficient earnings in covered employment and at least one day of covered employment in each of five years after 1936 and prior to reaching age 65.

ments. The major source of expansion of the program in subsequent decades was the aging-in of cohorts who had worked a sufficient length of time to be eligible for benefits. The growth of the program as it matured is shown in Figure 2, which shows the share of men receiving retirement benefits over time, separately by 5-year birth cohorts. Only about 5 percent of the 1871–75 birth cohort received monthly payments in 1940 at ages 65–69, rising to about 42 percent of those surviving to 1955 (ages 80–84). The 1896–1900 birth cohorts, born 25 years later, received Social Security at far higher rates: nearly 70 percent of men collected benefits in 1965, at ages 65–69, and 96 percent of those surviving to 1980 collected them then, at ages 80–84.

It will be central to our empirical analysis that there was also growth in eligibility due to expansion of the types of employment covered by Social Security at different times. The original act and the 1939 Amendments covered wage and salary workers only, excluding the self-employed. Wage and salary workers in certain sectors were also excluded from coverage: in particular, domestic workers in private homes, agricultural wage workers, employees of nonprofit organizations, and local, state, and federal government employees.[3] For types of employment covered in the original Act, coverage began on January 1, 1937.[4]

Amendments to the Social Security Act, particularly between 1950 and 1956, significantly expanded the definition of covered employment. In our analysis, our focus will be on wage and salary workers rather than the self-employed, so we first discuss those extensions. The first extension of coverage to additional categories of wage workers was under the 1950 Amendments. Under these amendments, coverage was made compulsory for regularly employed agricultural wage workers and domestic workers. The amendments also extended coverage on an elective basis to most employees of nonprofit organizations and to employees of state and local governments who were not under an existing retirement system. Importantly, the new additions to covered employment applied only to work performed in 1951 and later: coverage was not retroactive. Hence, a worker working solely in 1950-covered employment would not begin to accumulate quarters of coverage until 1951. Amendments in 1954 extended coverage to state and local employees already under a retirement system, still on an elective basis, and amendments in 1956 further liberalized the provisions under which state and local government employees could be covered. Although coverage was on an elective basis for nonprofits, take-up was fairly rapid: as of 1951, 56% of employees of

[3]See Appendix Section A.2.1 for a more detailed description of the initial exclusions and subsequent expansion. Nonprofit organizations were excluded due to concerns that it would jeopardize their tax-exempt status (McCamman, 1948), and state and local government employees were excluded because of perceived constitutional limitations on federal taxation of state and local governments (Clark et al., 2011). Other exclusions included railroad workers, who were covered under a separate system under the Railroad Retirement Act, and some other groups of smaller size, such as crew members of ships and fishermen.

[4]Work done at age 65 or older was not covered under the original act, but was covered under the 1939 Amendments, with an effective date of January 1, 1939. Hence, work during at least some part of 1937 and 1938 did not contribute to eligibility for workers in the 1873 and earlier birth cohorts.

nonprofit organizations were covered under Social Security, rising to 76% by 1960 (Subcommittee on Social Security, 1976, p. 27). For state and local governments, take-up was slightly slower than for nonprofits, but still rose to high levels: about 11 percent of all state and local government employees were covered by Social Security at the beginning of 1951 (the first year in which they could be), and by 1958 the majority of state and local government employees (54 percent) were covered by Social Security (Kestenbaum, 1982).

Coverage of self-employment was also important to the growth of Social Security, although it does not play a role in our empirical analysis. The 1950 Amendments extended coverage on a compulsory basis to the self-employed, except for certain professional groups and farmers. Self-employed farmers and farm managers were largely brought into the system in 1954, with a further expansion in 1956. Self-employed professionals who had not been covered in 1950 were mostly covered in either 1954 or 1956.

Under the 1939 Amendments, a worker who had worked a sufficient number of quarters in covered employment was eligible to receive monthly retirement benefits after reaching age 65, subject to an earnings test that withheld benefits in months in which earnings exceeded $15 (in 1939, this amount corresponded to about $230 in 2010 dollars). In addition, the spouse of a retired male worker received spouse's benefits (50 percent of the primary benefit) and the widow of a male worker received widow's benefits (75 percent of the primary benefit). The primary monthly amount was a function of the beneficiary's earnings in covered employment. Replacement rates for the median earner were around 20 percent through the 1940s, and 25 to 30 percent through the 1950s and most of the 1960s; for a low earner replacement rates were somewhat higher, around 25 to 30 percent in the 1940s and 40 to 50 percent in the 1950s and 1960s (SSA Office of the Chief Actuary, 2013, Supplementary Table V.C7). Tax rates for both the employer and employee were 1 percent beginning in 1937, the first year any employment was covered; these rose to 1.5 percent in 1950 and to 2 percent in 1954, and gradually rose in later years.

As has been frequently noted before (e.g., Moffitt, 1984; Schieber and Shoven, 1999; Gruber, 2013), the initial generation of beneficiaries paid little in taxes, received large benefits, and hence received large net transfers under Social Security. Famously, the first recipient of monthly benefits, Ida May Fuller, paid $24.75 in taxes and, by her death at age 100, had received $22,888.92 in benefits (all amounts nominal). As a more typical case, consider a married couple with a husband born in 1880 and a wife born in 1883. Suppose that the wife did not work, but the husband worked continuously in employment covered under the 1935 act, until he was 65, and in each year earned the average for all workers in covered employment. Also suppose that both husband and wife lived exactly their remaining life expectancy at age 65. The net present value at retirement of future benefits the couple received, minus taxes paid, is on the order of $84,000 in 2010 dollars. By

comparison, the median home value in 1940 was about $45,774 in 2010 dollars.[5]

2.2 Transfers and shared resources between the elderly and their adult children

For interpreting our empirical results, it is useful to discuss the typical nature of support between the elderly and their adult children prior to and in the early years of Social Security. We focus on discussion of support that children provided to their parents ("upstream" transfers), but discuss other types of relationships at the end of this section.

One of the major forms of support between adult children and their parents in the first half of the 20th century was co-residence. Figure 1 shows that in 1930, roughly 27 percent of women 65 and older, and about 12 percent of men 65 and older, lived with their children and were not the household head; co-residence rates of the elderly fell significantly in the subsequent decades.[6] Because we are interested in effects of Social Security on adult children, it is also useful to consider the share of adult children co-residing with parents, and hence what an adult's expectations might have been regarding their future living arrangements. Panel (a) of Figure 3 shows the share of native-born adults who were either household head or the spouse of the household head and who had one or more parents or parents-in-law in the household, tracing individual 10-year birth cohorts over their life cycles. For each birth cohort co-residence rises between ages 25–34 and 35–44 and then declines, presumably due to parents' mortality. For all cohorts shown here, this form of co-residence peaked at ages 35 to 44. For cohorts that reached their mid-20s by 1940, about 7 percent had a parent or parent-in-law in their household at ages 35–44, serving as a lower bound on the share of children who ever lived as household heads with their parents in the household. Note that since all children within a family are included in the base, with an average of between 4 and 5 children per family, the share of families in which children ever lived with their elderly parents was significantly higher. Consistent with evidence that Social Security led to reductions in co-residence of the elderly (e.g., McGarry and Schoeni, 2000), co-residence patterns changed significantly for cohorts reaching age 25 in 1950 or later, who lived with their parents at much lower rates.

Systematic data on other non-monetary types of help between adult children and their elderly parents, especially prior to Social Security, is more scarce than is information on co-residence. But evidence from the 1950s, when many elderly cohorts were still not eligible for Social Security, indicates that these types of support were widespread (Schorr, 1960). Indeed, they were similar

[5] See https://www2.census.gov/programs-surveys/decennial/tables/time-series/coh-values/values-unadj.txt, accessed 2021-03-18.)

[6] This series shows a decline prior to 1940 which Costa (1998) attributes in large part to rising income. The decline from 1930 to 1940 is most likely not due to Social Security retirement benefits, for which regular payments did not begin until 1940, but Old Age Assistance played an important role in the 1930–40 decline (Fetter, Lockwood and Mohnen, 2018). We focus on this measure of co-residence rather than a broader one in which parents are named the household head as a way to distinguish it from changes in children's propensity to leave their childhood home in young adulthood.

to the types of help by adult children prevalent in more recent periods (as described, e.g., by Horowitz, 1985). Several sociological studies of family connections of the elderly were carried out in the 1950s and early 1960s, some studies of a sample of households within a specific city, and some national. Notably, many of the major forms of non-monetary help that children provided their parents other than shelter required both time and spatial proximity. These types of help included physical care, shopping, and performing household tasks (Sussman, 1965). Help during illness, including long-term illness, was particularly common (Streib, 1958; Sharp and Axelrod, 1956).

Monetary transfers from children to their parents were present, but as emphasized by observers at the time (e.g. Schorr, 1960), connections within extended families did not primarily take this form. In a supplement to the 1952 Current Population Survey, Steiner and Dorfman (1957) found that for couples and unmarried men over 65, fewer than 1 percent received any regular monetary contributions from people outside their household; for unmarried women, less than 2 percent received regular monetary contributions as a major source of income. Irregular monetary contributions—such as paying a doctor or hospital bill in cases of serious illness—were probably more prevalent than regular monetary contributions. In a national survey of the elderly in 1957, Shanas (1961) found that there were contributions from children in about 20 percent of the cases in which the elderly had a recent serious illness. Nevertheless, Steiner and Dorfman (1957) regarded these irregular monetary contributions as being less important than co-residence as a form of support that adult children provided their parents. The reasons that monetary transfers were uncommon in comparison to other forms of help may have included a lack of markets for assistance with household tasks or shopping, or the ability to economize on both living costs and the time costs of physical or other care through co-residence. But it is also worth noting that in addition to their material needs, the elderly had strong desires for love and affection from their children. Respondents in Streib (1958), for example, reported that affection from their children was as or more important than financial help. Shanas (1961) found that in opinion surveys, many older people were concerned that asking for financial help from their children would threaten the affectionate character of their relationship with their children.

As they do now (Horowitz, 1985; Grigoryeva, 2017), daughters and sons differed in their connections with aging parents. For co-residence, asymmetry between daughters and sons widened in the early years of the Social Security era, then began to converge again. Panel (b) of Figure 3 shows time trends in the share of US-born men and women aged 25–54 living with their own parents. At the beginning of the 20th century, co-residence rates of older parents with daughters and sons were not substantially different, and in fact in 1900, co-residence with sons was slightly more common than with daughters. Over the middle decades of the 20th century, however, the share of sons coresiding with parents fell significantly faster than the corresponding rates for daughters; by 1960, daughters were much more likely to live with their parents than sons were. Interestingly, these

rates converged in 1990 before diverging again. There is, unfortunately, little systematic evidence from the early 20th century on children's provision of non-housing services, which are the type of support that survey data have found to be disproportionately provided by daughters in more recent periods (Horowitz, 1985). But in the 1957 National Survey of the Aged, Shanas (1961) found that approximately 60 percent of people 65 and older named a son or daughter as the person they would ask for help during an extended illness, and that they were about twice as likely to name a daughter as to name a son. Given the similarity to evidence from recent periods, it seems plausible to assume that the greater reliance on daughters than on sons for these types of care was also the case in the 1940s and earlier.

Downstream transfers, or more simply sharing of resources, were also prevalent. There was, for example, a significant degree of co-residence of adult children with older parents as the household heads, and regardless of headship status, co-residence may have been a matter not of transfers in one direction or another, but rather of mutual economic benefit (Schorr, 1960). Older parents also provided a substantial amount of non-housing services, such as help caring for young children or help during illness (Schorr, 1960; Sussman, 1965). Finally, there is more evidence for significant downstream financial transfers than for upstream financial transfers; several sources comment that the balance of financial transfers tended to be downstream (Sussman, 1953; Streib, 1958; Schorr, 1960; Sussman, 1965). It should be noted, however, that much of the evidence for significant downstream financial transfers is at relatively young ages, when parents are younger than 65 or not much older, and in the early years of children's marriage.

3 Data

3.1 Linked Numident-census data

Our primary linked sample relies on two data sources: the 1920 and 1930 full count US Censuses, and the public version of the Social Security Numerical Identification (Numident) File released by the National Archives and Records Administration (NARA). Contained in the Numident are Social Security death records for 49 million individuals. Our primary goal is to link individuals in death records, in which we can observe the ZIP code where the individuals last lived, to their fathers in the 1930 Census, which allows us to predict their fathers' future Social Security eligibility based on their age and employment information in 1930 (five years prior to the passage of the Social Security Act). Because death records only contain basic information on individuals (full name, date of birth, and place of birth), the basis of our linked sample is a second set of Numident records: Social Security application (form SS-5) records for 40 million individuals. The majority of SS-5 records are immediately linkable to death records via Social Security Numbers (SSNs), and have the advantage that they additionally list individuals' parents' full names. This allows us

to link individuals and their parents as one single family unit to households in the 1930 Census. To address the fact that some children may not be living with their parents anymore in 1930, we augment our sample by combining analogous linkages to the 1920 Census and linkages between households in the 1920 and 1930 Censuses.

We highlight two key advantages of our linked dataset relative to existing longitudinal datasets based on historical records. Typically, these datasets link individuals across records solely based their own information (e.g., name, year of birth, place of birth), and tend to focus on men, since women have historically changed their surname at marriage. In contrast, our ability to observe parent names in both SS-5 and census records allows for significantly more disambiguation, increasing the number and quality of links we are able to make. In addition, our dataset covers sons and daughters at similar rates. This is because men and women are present in SS-5 records at similar rates, and with information on parents' names, linking individuals to their childhood household is gender-neutral.[7]

Appendix Figure A.1 provides an overview of how we construct our linked sample. We start by linking siblings in SS-5 records together based on parent names. We then link these reconstituted families to households in the 1920 and 1930 Censuses based on parent names and information on children (first and middle name, age, and place of birth). Next, we link households in the 1930 Census to households in the 1920 Census using the same information as the previous links, except that we can additionally exploit parents' age and place of birth. Lastly, we restrict our sample to individuals in SS-5 and death records that can be linked to both the 1920 and 1930 Censuses via any combination of the linkages just described. Although we are mainly interested in observing fathers in 1930, observing individuals' household structure in 1920 is important because it allows us to better approximate their true family size, a key variable in our analysis.[8]

Each of the linkages described above involves linking records primarily based on names. In this paper, we adopt the supervised machine learning (ML) approach to historical record linking from the Longitudinal, Intergenerational Family Electronic Micro-Database (LIFE-M) Project (Bailey et al., 2021).[9] Like other supervised ML approaches, this method involves making manual linking decisions for a random sample of records, which are then used to train an algorithm to make

[7] Although women participated in the labor market at lower rates than men historically, SS-5 records contain the spouses of primary earners as spouses were eligible for benefits after the 1939 Amendment to the Social Security Act. Note that for married women, whose reported surname in SS-5 records is their husband's surname, the ability to observe their surname at birth via their father's surname is what enables us to link them to their childhood household.

[8] The censuses only provide information on family size at one specific point in time, conditional on co-residence. Alternatively, we could infer family size from reconstituted families in SS-5 records, but they are often incomplete as only a subset of siblings appear in the SS-5 data.

[9] Other supervised ML approaches to historical record linking include the IPUMS Linked Representative Samples (Ruggles, 2006), Feigenbaum (2016), and the IPUMS Multigenerational Longitudinal Panel (Helgertz et al., 2020). A common alternative approach, though not the one we take here, is to use automated linking algorithms (Ferrie, 1996; Abramitzky et al., 2020, 2021).

"similar" linking decisions for millions of records. A key feature of the LIFE-M training process is that human trainers examine an entire set of potential candidates ("potentials") for each record we want to link ("primary"), and decide whether or not to make a link and which link to make. Human input is useful as there are often multiple plausible candidates due to name commonality, and deciding which candidate to choose or concluding there is too much ambiguity to do so is not something that can easily be automated using deterministic rules. One valid concern is that the quality of the model links will only be as good as the quality of the human links (Abramitzky et al., 2021). We mitigate this in two ways: (1) by instructing trainers to only make links when they are very confident, and (2) by aggregating decisions from multiple trainers, as links that are independently chosen by different trainers are more likely to be objectively good links.[10]

The resulting hand-linked data is used to discipline a machine learning algorithm to mimic human decisions. We use the two-stage model from Bailey et al. (2021), which is specifically designed for one-to-one linking.[11] The output of the model is an estimated match probability for every primary-potential pair. The model-based links are then determined using a threshold rule: potentials that have a match probability exceeding the threshold are called links. A key decision is how to set this threshold. In general, the lower the threshold the higher match rate but the lower the agreement rate with trainers ("precision rate"). Conversely, the higher the threshold the lower the match rate but the higher the agreement rate, creating a natural trade-off between making more links and fewer false links. Training data allows us to control this trade-off, and we select the threshold associated with a 97 percent precision rate for all our linkages. Using the trained model and the chosen threshold, linking can be scaled to the full sample of records. Details for each specific link, including how we generate sets of potential candidates, what information is shown to trainers, and how well our models approximate trainer decisions, are provided in Appendix A.1.

After combining links between SS-5 records, death records and the 1920 and 1930 Censuses, we restrict the sample to our population of interest: children whose fathers were born between 1875 and 1888 and covered by Social Security in 1935 or 1950 based their employment information in

[10]Following the LIFE-M training process, trainers first complete a training module and are given feedback so that they can familiarize themselves with each specific linkage. Two trainers then make linking decisions for the full random sample of primaries, and their decisions are compared. Disagreement cases are shown to three additional trainers (mixed in with a random sample of other cases). Only links chosen by the two original trainers, or links chosen by four out of five trainers are considered true links for the purpose of training the model.

[11]In the first stage, the probability that there exists a link in the set of potential candidates is modeled using a random forest model (Friedman et al., 2009). The random forest learns to label entire sets as containing a link or not based on set-level features, and is fitted on the training data collapsed at the primary level. In the second stage, the probability that any link in the set is the true link conditional on there being a link in the set is modeled using a logit-style discrete choice model. The explanatory variables in the logit model are pair-level features between each primary and potential, and the associated parameters are estimated via maximum likelihood using the subset of primaries for which a link was found. The stage 1 and stage 2 probabilities are generated using 10-fold cross-validation to minimize overfitting, and multiplied to obtain the unconditional match probability for each primary-potential pair.

1930, excluding self-employed workers.[12] The resulting linked sample contains around 1.1 million children born between 1895 and 1930 with non-missing ZIP at death information, belonging to around 650,000 unique families. Appendix Figure A.2 plots the cohort distribution of fathers and children in the population and in our linked sample. The coverage rate of fathers is around 22 percent, while the coverage rate of their children born by 1930 is slightly lower at around 13 percent.[13]

Several factors contribute to incomplete coverage. First, only a subset of children appear in the Numident, either because they never applied for a Social Security Number or due to restrictions placed on the public version of these files. In particular, the Numident file released by NARA only includes individuals who died prior to 2007 and whose deaths were not state-reported. In practice, this file has near-universal coverage of deaths that occurred between 1988 and 2007 (Goldstein and Breen, 2020), and relatively low coverage of deaths that occurred prior to 1988. Second, ZIP at death information is missing for around 27 percent of death records. Third, human trainers are only able to confidently link a subset of records, and in turn the model is only able to replicate a subset of the links made by human trainers at a 3 percent error rate. Lastly, around 6 to 8 percent of raw occupation and industry strings digitized from 1930 Census forms have yet to be classified.

Incomplete coverage raises two types of selection concerns: (1) which families appear in our linked sample, and (2) which children appear in our linked sample. Appendix Table A.1 assesses the representativeness of our linked sample in terms of 1930 household characteristics. Overall, our sample is relatively balanced, though a few patterns are worth noting. First, as is common in historical linked samples, Black and low-socioeconomic status families (e.g., as measured by father literacy or home ownership) are underrepresented as they tend to have more common names and report their names less consistently across records. Second, because the population includes men that never had any children and because observing multiple children improves our chances of making a link, our linked sample is naturally skewed towards families with a greater number of children in the household in 1930. To account for these differences, we re-weight families in our linked sample using inverse propensity score weights (IPWs) generated using the procedure described in Bailey et al. (2020).[14] By design, the resulting weighted means in column 3 are very

[12] We choose the cutoff points of 1875 and 1888 for two main reasons. First, to identify a father-child relationship between two people, we must observe children with their father in the same household in 1920 or 1930. Including men born earlier will increase the number of children who have left the household by 1920, and including men born later will tend to increase the number of children not yet born by 1930. Second, as shall be evident in Section 4, the birth cohorts from 1875 through 1888 cover close to the full range of differences in our key policy variable between 1935- and 1950-covered workers.

[13] We cannot compute the exact coverage rate of children since our population of interest conditions on fathers' employment status in 1930 (and not all children live with their fathers in 1930). Instead, 13 percent is the coverage rate of children whose fathers were born in 1875–1888 regardless of their 1930-based Social Security coverage status, which we compute by pooling the number of children aged under 10 (or 15) across the 1910–1940 Censuses. Our coverage rate is an underestimate due to the fact that some of those children never reached adulthood.

[14] The weighting variables are fixed effects for father year of birth, race, place of birth, and literacy status, as well as

close to the population means in column 1.

Intuitively, selection of families into our linked sample is mild because most families in our population of interest had at least one child who ever applied for an SSN and died between 1988 and 2007. However, it is clear that our sample contains the subset of children who satisfy these criteria. In particular, as was evident in Appendix Figure A.2, our linked sample is skewed towards later cohorts of children as they were more likely to survive until 1988. This reflects the selection of cohorts into SS-5/death records, prior to linking (see Appendix Figure A.3). On the other hand, selection in terms of place of birth or place of death is less of a concern given that SS-5/death records have excellent geographic coverage (see Appendix Figure A.4). Selection in terms of which families or which children appear in our linked sample matters only to the extent that it is correlated with the Social Security coverage variation we exploit. In Section 4.4 we show that this is not the case, alleviating the concern that selection would affect the conclusions of our analysis.

The Numident also includes a third set of files: Social Security claims records for 25 million individuals. We build a separate dataset linking men born in 1875–1888 in the claims records to themselves in the 1930 Census. We use this sample in the analysis to assess whether our prediction of fathers' eligibility for Social Security based on their 1930 employment information is predictive of their actual claiming behavior, as measured by their probability of ever appearing in the claims data or the age at which they claimed Social Security benefits conditional on appearing in the claims data. These linkages are made on the basis of name, age, place of birth, and place of residence information using the same supervised ML approach described earlier (see Appendix A.1 for more details). Ultimately, we are able to link around 360,000 men covered by Social Security in 1935 or 1950 at a 97 percent precision rate, or around 10 percent of the population. We re-weight this sample using the same procedure used to re-weight our main sample (see Appendix Table A.2 for summary statistics).

3.2 Survey of Occupational Changes in a Generation

While our core analysis is based on the Numident-census linked sample described in the previous section, we also present complementary survey-based evidence. Specifically, we exploit the 1973 wave of the Survey of Occupational Changes in a Generation. The OCG was a supplement to the 1973 Current Population Survey (CPS) that focused on civilian males aged 20–65 and collected detailed information on the respondents themselves as well as their parents and spouses. Crucially, it contains retrospective information on the respondent's father's occupation, industry, and self-employment status when they were aged 16, as well as the father's year of birth. Combining this

fixed effects for household home ownership, radio ownership, farm status, and urban status. Other examples of studies using re-weighting techniques to adjust for selection include Pérez (2017) and Zimran (2019).

information enables us to predict fathers' likely Social Security eligibility age.[15]

A unique feature of the 1973 CPS is that it was linked to tax records from the Internal Revenue Service (IRS) and earnings histories from the Social Security Administration (SSA) in a joint effort with the U.S. Census Bureau. In Appendix A.1, we describe how we link together the 1973 OCG file (Blau et al., 1994), the 1973 CPS file (United States Bureau of the Census, 1992), the Social Security Records: Exact Match Data file (Social Security Administration, 2005), and the 1937–1975 Social Security Longitudinal Earnings file (Social Security Administration, 2008), despite the absence of unique identifiers.[16] After restricting the sample to sons born in 1907–1930 whose fathers were born in 1875–1894 and covered by Social Security in 1935 or 1950 based on their retrospective employment information (excluding self-employed workers), we are left with 3,020 observations.[17] In addition to containing detailed income information, a key advantage of the OCG sample is that it is nationally representative.

4 Empirical Strategy

4.1 Parameterizing variation in Social Security eligibility

In our empirical analysis, we map an individual's birth cohort and type of employment in 1930 into that person's likely eligibility for Social Security. Doing so addresses two issues. First, we do not directly observe taxes to or transfers from Social Security. Second, we want a measure of likely Social Security eligibility that is not a function of behavioral responses to Social Security. To parameterize variation in Social Security, we use the 1935 Act and the 1939 and 1950 Amendments to calculate the earliest possible age at which an individual from a given birth cohort and with a given employment type could stop working and be permanently fully insured (and hence eligible to receive benefits upon reaching age 65), assuming that workers do not change the type of their employment to one that was covered earlier.[18] Although some workers surely changed from 1950-covered employment to 1935-covered employment, we will be able to assess the degree to which

[15] A first wave of the OCG was carried out in 1962, but this survey did not collect information on fathers' year of birth, which is needed to determine their likely Social Security eligibility age.

[16] In short, we exploit the fact that these files contain common variables that uniquely identify nearly all individuals we are interested in.

[17] 1907 is the earliest birth cohort in the OCG data (aged 65 at the time of the 1973 CPS), while 1930 is chosen as an upper bound for consistency with our Numident-census linked sample and because we are interested in children's outcomes in adulthood. Relative to the Numident-census linked sample, we also expand the set of father cohorts to increase sample sizes.

[18] "Permanently" fully insured means that a worker could stop working in covered employment and still be fully insured—eligible to receive benefits—upon reaching age 65. The distinction arises from the fact that a worker can be "fully insured" in case of death but lose fully insured status if he or she ceased to work in covered employment. In the discussion that follows, we use the terms "permanently fully insured" or "permanently insured." In our analysis, the distinction is of modest importance because the cohorts we focus on were largely bound by the 6-quarter minimum.

employment in 1930 is predictive of later claiming of Social Security.

Figure 4, panel (a), illustrates the minimum age at which a worker from a given birth cohort would become permanently fully insured under the 1950 Amendments, under the assumption that those in 1950-covered employment do not engage in any 1935-covered employment. Following this assumption, we will sometimes refer to workers as "1935-covered workers" or "1950-covered workers" even though workers could, of course, switch between 1935-covered and 1950-covered work. For hypothetical workers from a given birth cohort in different types of employment, the required number of quarters of coverage is the same; the difference between the two in the earliest age of reaching permanently insured status lies in the 14-year delay between 1935-covered work contributing to eligibility (beginning in 1937, except for those 65 or older in 1937 and 1938) and 1950-covered work contributing to eligibility (beginning in 1951). Compared to a worker in 1950-covered employment, a worker in 1935-covered employment is more likely to have worked long enough to ever be eligible, due to exit from the labor force at older ages. A worker in 1935-covered employment will also be permanently insured at an earlier age, conditional on ever being eligible. Although the difference in age at reaching permanently insured status is largely constant across birth cohorts, the effect of this difference on the likelihood of ever being eligible is heterogeneous given latent retirement behavior. Since exit from the labor force at older ages (say, 65 to 80) is greater than exit at younger ages (say, 50 to 65), the difference between 1935-covered workers and 1950-covered workers in ever being eligible tends to be greater for older cohorts, at least over some range. Panel (b) of Figure 4 shows the earliest age at which a beneficiary could begin collecting monthly benefits, which is the primary parameterization of "exposure" to Social Security that we use in our analysis. As is also evident from panel (a), a worker who remains in 1950-covered employment reaches eligibility at an older age. But since no workers received monthly benefits until age 65, the difference in the earliest possible age of receiving monthly benefits is smaller for later birth cohorts.[19]

4.2 Comparing similar workers with differing Social Security eligibility

We get empirical traction on this parameterization of likely Social Security eligibility by observing an individual's age (and hence birth cohort) and employment information in the 1930 Census. For all workers, the Census recorded an occupation, industry, and class of worker (e.g. self-employed, working for wages or salary, or unpaid). We classify all combinations of these three variables into a year when this type of employment was covered by Social Security. Full details are provided in Appendix Section A.2.1, and Appendix Table A.3 shows that in the 1940 Census, our coding is highly predictive of workers reporting that they had a Social Security Number. A priori, one

[19]Men were eligible to claim at ages 62 to 64 starting in 1961, which is relevant only for cohorts born in 1897 and later.

would expect that the types of workers who were covered in 1935 would systematically differ from those covered in 1950; the self-employed, domestic workers, and agricultural wage workers would likely have different underlying characteristics than wage and salary workers in manufacturing, for example. To the extent that these differences are level differences, constant across birth cohorts, one could net these differences out in a specification of the form

$$y_i = \alpha_{b(i)} + \beta_{c(i)} + \gamma \cdot (\text{min. age of Social Security eligibility})_i + \Lambda' \mathbf{x}_i + \varepsilon_i, \quad (1)$$

regressing the outcome y for family i on fixed effects for father's birth cohort b and father's coverage category c, with γ the coefficient of interest and $\mathbf{x}_i$ a vector of controls.

In practice, when making all possible comparisons between families of workers covered in the different years, the average characteristics of the groups in 1930 appear to be different. Hence, our primary specification is based on a finer comparison. In particular, the exclusion of certain industries in the original act generates variation that allows us to exclude the self-employed and to make comparisons of wage and salary workers with the same occupation, but working in different industries. As noted earlier, employees of nonprofit organizations, state and local governments, and private households were not covered under the 1935 Act or 1939 Amendments, even though work in the same occupation in other sectors was covered by Social Security. For example, work as a janitor for a manufacturing firm would have been covered in 1935, but work as a janitor in a school, a nonprofit hospital, or a church would not have been covered until at least 1950. Differing treatment of similar workers did not escape the notice of the Social Security Board: Arthur J. Altmeyer, its Chairman, noted in 1944 that many employees of nonprofit institutions were not covered, despite the fact that their "skills, tasks, and earnings... do not usually distinguish them from comparable employees in commerce or industry" (Altmeyer, 1944).

To compare workers in the same occupation but with coverage status differing due to industry of employment, we exclude self-employed workers from our sample and estimate a specification of the form

$$y_i = \alpha_{b(i)} + \beta_{c(i)} + \gamma \cdot (\text{min. age of Social Security eligibility})_i + \delta_{b(i)o(i)} + \Lambda' \mathbf{x}_i + \varepsilon_i, \quad (2)$$

where, relative to equation (1), we include fixed effects for the interaction of father's birth cohort and father's 1930 3-digit occupation (subscripted by o) in order to focus on variation in minimum age of eligibility arising from work in different industries. Equation (2) serves as our preferred empirical specification.[20] As described in detail in Appendix Section A.2.1, we use the industry

[20]The four coverage categories are: covered in 1935, covered in 1950 (domestic work), covered in 1950 (nonprofit sector), and covered in 1950 (state and local government). Our baseline controls are fixed effects for father race and state/country of birth. Observations are weighted using the inverse probability weights from Appendix Table A.1.

classification of the Census to identify employees who are less likely than others to have been covered by the original 1935 Social Security Act. Work for private households was identified as a separate industry category. Nonprofit organizations were concentrated most heavily in hospitals, educational institutions, and welfare and religious organizations, in addition to some other nonprofit membership organizations. Hospitals and educational services were also important areas of state and local government employment, in addition to more general state and local public administration. Hence, we treat wage and salary workers in these industries as not covered until 1950.[21] Note that while most employment in hospitals and educational institutions was in nonprofit organizations or in state and local governments, there were a handful of for-profit organizations in these sectors, as we discuss in Appendix Sections A.2.1 and A.3. Any misclassification of this sort should tend to attenuate our results.

The occupations that form the basis of our comparison are those that were present in both for-profit and either nonprofit or public enterprises. Appendix Table A.4 shows the 20 largest occupation categories in the population and our linked samples in which there is non-trivial variation in Social Security coverage status (defined here as the smaller group comprising at least 1 percent of that occupation). Among the largest categories are janitors, truck drivers, and clerical workers.[22]

For interpretation of our estimates, it is useful to note that some workers we classify as not covered by Social Security until at least 1950—state and local government workers in particular—may have been covered by existing retirement systems. We discuss existing retirement systems for government and nonprofit organizations in detail in Appendix Section A.3, and here provide a few notable facts. At the beginning of 1942, when the Census Bureau carried out the first systematic national survey, 46 percent of all state and local government employees were covered by a retirement system (United States Bureau of the Census, 1943, Table 2). However, the types of employees most often covered by these retirement systems, such as police, firefighters, and teachers, are either excluded from our sample or have little effect on the estimates because these occupations do not appear in private industry. Government employees that form the primary basis for our comparisons—such as custodial and maintenance employees—were less likely to be covered under existing retirement systems. As for nonprofit organizations, about a third of private, nonprofit institutions of higher education had retirement plans in the 1940s, but the maintenance and non-

[21] As described in Section 2.1, nonprofit employees and state and local workers not covered under an existing retirement system were covered on an elective basis starting in 1950, while state and local workers already covered under a retirement system were not covered by Social Security until 1954. In the Census we do not observe whether or not workers are covered by a retirement system, nor can we generally distinguish between employees of nonprofit organizations and employees of state or local governments. Hence, we treat all such workers as covered in 1950.

[22] Note that many of the wage workers covered in 1950 for whom it is less clear that there would be similar 1935-covered workers do not play a significant role in our comparisons. In particular, most agricultural wage workers and many domestic workers are in their own occupation codes in the census classification. They are included in our sample to help identify covariates, but with no variation in their coverage status, do not contribute significantly to our estimates of interest.

professional employees who are our focus were generally not eligible (McCamman, 1948). Only about 8 percent of nonprofit hospitals had pension plans at the beginning of the 1940s (Hayhow, 1940). Finally, it is also useful to note that the benefits under retirement systems for both public and nonprofit employment typically differed from Social Security in important ways. Most relevant to our analysis, these plans tended to have either no provision or less generous provisions for widow's benefits (McCamman, 1953). Widows are a particularly relevant case for our study, since co-residence and other types of family old-age support were especially common among widows.

4.3 Predicted Social Security eligibility and subsequent claiming behavior

The discussion above suggests two testable predictions. First, workers whose pre-Social Security employment would be covered in 1935 should be more likely to ever claim Social Security than workers whose employment would only be covered in 1950, and conditional on claiming they should claim at earlier ages. Second, these differences should be smaller for later birth cohorts than earlier birth cohorts, at least over some range. We test these predictions using our sample of Social Security claims records linked to the 1930 Census. For individuals found in the claims data, we can observe the age at which they claimed Social Security benefits. While we cannot distinguish between individuals who never claimed Social Security and individuals who did but were not linked to the claims data (or simply do not appear in it), ever claiming should be positively correlated with the probability of being linked to a claims record.

We begin with a descriptive comparison. First, in the 1930 Census, 1935-covered men born in 1875 are three times more likely than 1950-covered men to be linked to a claims record. This figure drops to one and a half for men born in 1888. Next, focusing on the subset of men linked to a claims record, Figure 5 plots the distribution of age at claim by coverage status, separately for men born in 1875 and 1888. For the 1875 birth cohort, it is evident that 1935-covered workers claimed benefits earlier (3.6 years on average). In contrast, the age at claim distributions are much more similar for the 1888 birth cohort (average difference of 2.2 years), which reflects the fact that even workers who never switched out of 1950-covered employment would have satisfied the eligibility requirement by the time they reached age 65 in 1953.

Our parameterization of Social Security eligibility is also predictive of ever claiming Social Security and age at claim. Table 1 reports estimates from equation (2), which exploits variation in eligibility age arising from workers with the same occupation but working in different industries. In column 1, the sample is the 1930 population of men born in 1875–1888 and covered in 1935 or 1950 (excluding the self-employed), and the dependent variable is an indicator for being linked to a claims records. In column 2, we focus on the subset of men linked to a claims record and the dependent variable is the age at which they claimed Social Security benefits. Column 1 shows that eligibility one year earlier is associated with a 0.33 percentage point increase in the likelihood of

being linked to a claim for monthly benefits, relative to a baseline mean of 10 percent. Column 2 shows that, conditional on being linked to a claims record, eligibility one year earlier is associated with claiming monthly benefits 0.12 years earlier. Note that we do not formally treat these estimates as a first stage—and in particular do not use them to scale our reduced-form estimates—because neither fully captures the way in which a father's earlier eligibility for Social Security may influence his children's outcomes. Nevertheless, they do establish that there is a strong relationship between our parameterization of Social Security eligibility and subsequent claiming behavior, even when restricting comparisons to fine occupational categories.

4.4 Selection into the linked sample

Due to the coverage of Numident records and the fact that we can only link a selected subset of them to the censuses, the families and children that appear in our linked sample are a selected subset of the population. Even though we re-weight families in our sample to resemble the population in terms of observable characteristics, non-random selection could potentially affect the internal validity of our estimates if it is correlated with the variation we exploit. To assess the extent to which selection is a concern, we examine all households in our population of interest in 1930—those headed by a man born in 1875–1888 and covered by Social Security in 1935 or 1950 based on his employment information (excluding self-employed workers)—and test whether appearing in our linked sample is correlated with our Social Security eligibility variation. The results are reported in Table 2. Reassuringly, we find no evidence that fathers' minimum age of Social Security eligibility is correlated with the likelihood of the household appearing in our linked sample, the likelihood of any son or daughter appearing in our linked sample, or the number of sons or daughters appearing in our linked sample.[23]

4.5 Placebo results in the 1930 Census

A key check on whether our empirical approach allows for valid causal comparisons is to test whether our variation is correlated with household characteristics in 1930, five years prior to the introduction of Social Security. The set of characteristics we examine are home ownership, house value, radio ownership, urban and farm status, number of children present in the household, father and mother literacy status, and mother labor force participation status. Table 3 reports estimates based on two approaches, first testing for differences in the full population of households in panel (a), and then testing for differences within the subset of households contained in our linked sample in panel (b). The results confirm that conditional on making comparisons between households

[23] In Appendix Table A.5, we also find no evidence of selection of children from specific cohort ranges, with the exception of daughters born in 1911–1920. As we will show, our main findings are driven by sons so we do not view this as a major concern.

with fathers in the same occupation, there were no significant pre-existing differences between households whose fathers were predicted to be eligible for Social Security earlier versus later. With the exception of farm status in the linked sample, which is statistically significant but small in magnitude, all coefficients are close to zero and statistically insignificant. To further confirm that observable differences across households are not driving our results, we will show that our results are robust to controlling for 1930 household characteristics.

5 Results

5.1 Characteristics of late-life location for sons and daughters

Having established that the variation we exploit is correlated with observed Social Security claiming patterns, but not with pre-existing family characteristics or selection into the linked sample, we now explore how fathers' Social Security eligibility affected the long-run outcomes of their children. We proxy for children's long-run socio-economic status using the "quality" of the ZIP code they last resided in prior to their death. A large literature suggests that there is a strong relationship between neighborhood quality and individual-level socioeconomic and health outcomes (see Chyn and Katz, 2021, for a recent review). In some sense, the neighborhood in which children died can be viewed as cumulative measure of how well they fared over their lifetime. In practice, we measure neighborhood quality using ZIP-level mean adjusted gross income (AGI) in 2001 or median house values in 2000, though we will show that similar results are obtained using other commonly-used metrics such as poverty rates.[24]

Throughout the analysis, we allow for heterogeneity along two key dimensions. First, we allow for heterogeneity by the number of children in the family. It is likely that the burden of caring for aging parents is greater for the average child from a small family than the average child from a large family, simply because there are fewer children to support the parents. Second, we examine impacts on sons and daughters separately. As described in Section 2.2, sons and daughters differed in the types of support that they provided their parents. In particular, while sons and daughters coresided with their parents at broadly similar rates early in the 20th century, the evidence suggests that there were more substantial differences in their propensity to provide support that was not easily purchased on the market, such as personal care. This difference suggests that Social Security might displace transfers from sons to a greater degree than transfers from daughters, and hence that sons may exhibit a larger out-migration response to their parents' receipt of Social Security.

Figure 6 plots the OLS estimates from equation (2), where the dependent variable is the ZIP

[24]Data on ZIP-level mean AGI comes from the Statistics of Income Division of the IRS, while data on ZIP-level median house values comes from the National Historical Geographic Information System (NHGIS, Manson et al. 2020). Measuring average ZIP income using median household income instead yields very similar results.

percentile rank within the national distribution of ZIP codes in terms of mean AGI (panel a) or median house values (panel b), averaged across all children observed in our linked sample for each family. We average outcomes across all children because father's Social Security eligibility is a family-level treatment. In each graph, the left panel displays the coefficient on father's minimum age of Social Security eligibility imposing a common effect across all family sizes, while the right panel displays coefficients from a separate regression where minimum eligibility age is interacted with a set of family size fixed effects.[25] We run separate regressions looking at outcomes for all children, sons only, or daughters only (indicated in the legend).[26]

The negative coefficients in both panels of Figure 6 imply that a *lower* minimum age of Social Security eligibility—in other words, fathers' greater likelihood of eligibility or earlier eligibility—has a *positive* impact on children's end-of-life neighborhood quality, whether measured by income or house prices. Consistent with the notion that a shift from family-based to government-based old-age support has greater a impact on children from small families, the effects are larger in absolute terms and more often statistically significant for children from smaller families. Consistent with the prediction of differential effects for sons and daughters, we find that the effects are entirely concentrated among sons, with little to no effects among daughters. In terms of magnitudes, a ten-year difference in fathers' minimum age of Social Security eligibility—roughly the difference between 1935- and 1950-covered fathers from the 1877 birth cohort—is associated with a 2 percentile difference in the income or house price rank of sons' ZIP at death. Appendix Figure A.5, in which ZIP ranks are replaced with logs of ZIP outcomes, implies that these effects correspond to a 2.3 percent increase in ZIP-level AGI and a 2.8 percent increase in ZIP-level house values. These effects are roughly 50 percent larger in magnitude for sons from families with fewer than five children.

Our findings are robust to alternative ways of ranking ZIP codes. Appendix Figure A.6 shows results for sons and daughters based on four different metrics, all drawn from NHGIS: the share of college-educated individuals in 2000, the share of individuals under the poverty line in 1999, the share of households receiving public assistance in 1999, and the share of single-parent families in 2000 (the signs of the coefficients for the last three outcomes have been inverted so that a negative coefficient is "desirable" for all outcomes). Across all outcomes we find similar results: fathers' greater likelihood of eligibility or earlier eligibility has a positive impact on sons' end-of-life neighborhood, particularly for sons from smaller families. Appendix Figures A.7 and A.8

[25] We approximate true family size by aggregating all children present in the household in 1920 and the subset of children born after 1920 present in the household in 1930. Note that this approach does not capture children we never observe, and does not account for the possibility that some children born before 1920 may only appear in 1930 or the fact that age can be mismeasured in the census.

[26] Note that the underlying set of families implicitly changes when looking at sons or daughters, as not all families have at least one daughter and one son, and even if they do we might only observe a subset of them in our sample.

show that our results are also robust to controlling for 1930 household characteristics, adding 1930 county fixed effects, and not re-weighting observations using inverse probability weights. The county fixed effects in particular alleviate the concern that our findings might reflect differences in characteristics of 1930 locations correlated with Social Security eligibility.

As we have emphasized, a possible channel through which parents' receipt of Social Security may affect their children is through a greater likelihood of migrating. For example, to the extent that Social Security displaces upstream transfers that require physical proximity, it may induce some children to switch from co-residing with their parents to living independently, possibly in a different location, and other children to switch from independent living in the same location to moving to a different location. To assess the degree to which parents' receipt of Social Security affected children's long-run migration patterns, Figure 7 shows the impact of fathers' Social Security eligibility on the share of children who, at the end of their lives, were living outside their parents' 1930 place of residence. Parents' location in 1930, when they are in their 40s and 50s, should be viewed as a proxy for where the parents spent their last remaining years, and therefore where children would have had to live in order to take care of them. Although we only observe children's last known location, it gives us a sense of how many children migrated away from their parents' home (or did not return, if they had already left by 1930). We distinguish between two types of migration events: dying outside parents' 1930 county of residence, and dying outside parents' 1930 commuting zone of residence. The latter captures the extent to which children migrated to new local labor markets, which may have offered them better job opportunities.[27]

The negative coefficients in Figure 7 imply that fathers' eligibility for Social Security led to a greater share of children settling outside their parents' 1930 county and 1930 CZ of residence. Among sons, a ten-year difference in father's earliest age of Social Security eligibility is associated with a 3 percentage point greater likelihood of dying in a different county and a 2 percentage point greater likelihood of dying in a different CZ, which roughly correspond to 5 percent of the baseline means in our linked sample (see Appendix Table A.6). Similar to our findings for sons' ZIP-level characteristics, there is a strong family size gradient: the effects among sons from smaller families are around twice as large as the baseline effect.[28]

While there is some evidence that daughters from small families lived outside their parents' 1930 county of residence, migration responses for daughters are, on the whole, more muted than for sons. In particular, we find no evidence that the average daughter was more or less likely to

[27] Commuting zones are the standard way to define modern local labor markets. More specifically, they are clusters of counties characterized by strong commuting ties within CZs and weak commuting ties across CZs, based on commuting patterns in the 1990 Census (Tolbert and Sizer, 1996). Note that local labor markets were likely smaller than CZs at the time given that transportation options were much more limited.

[28] Appendix Figure A.9 shows that we get similar results if we look at the share of children who died outside their parents' 1930 state or Census region of residence. Appendix Figures A.10 and A.11 show that our county and CZ migration results are also robust to alternative specifications.

die in her parents' 1930 location. And except for only-child daughters, we find no evidence that father's Social Security eligibility made daughters more likely to settle in a different commuting zone.

5.2 Labor market outcomes for sons

The findings in previous section imply that sons whose fathers were more likely to be eligible for Social Security, or eligible at an earlier age, ended up in higher-SES ZIP codes. To what extent does this reflect real improvements in individual-level outcomes? In this section we provide some suggestive evidence based on the 1973 OCG survey, in which we can observe earnings and occupational standing in adulthood for a small but nationally representative sample of sons.

Figure 8 displays OLS estimates from equation (2). Two minor differences relative to the empirical setup for Numident-based results are that the units of observation are individual sons (as there are unlikely to be many siblings in our OCG sample, and we could not identify any sibling relationships that are present), and that we only present coefficients imposing a common effect across all family sizes due to small samples. Furthermore, the baseline controls are fixed effects for sons' race and state/country of birth (rather than fathers), and we additionally include son cohort fixed effects to account for life cycle differences in earnings and occupational standing. Observations are weighted using OCG sampling weights. Note that OCG respondents were asked about their father's employment information when they were 16, which might have endogenously responded to the introduction of Social Security in 1935 for fathers of sons born after 1919. We therefore show separate results for sons born in 1907–1919 to ensure that our findings are not driven by this endogeneity issue.

In the left panel, we examine the following income measures: cumulative Social Security earnings for four different time periods (1937–1950, 1951–1960, 1961–1970, 1971–1975), AGI in 1972, and IRS salary and wage in 1972 (see Appendix Table A.7 for summary statistics).[29] The negative coefficients, though imprecisely estimated, broadly suggest that sons whose fathers were more likely to be eligible for Social Security, or were eligible earlier, had higher Social Security earnings between 1937 and 1950, and higher income in 1972. The coefficients for Social Security earnings over the period 1951–1960 are negative but not statistically significant, and there are no meaningful differences in Social Security earnings after 1960. Although income measures are either missing or zero for 10 to 20 percent of sons (depending on the measure), Appendix Figure A.12 shows that the results are not driven by a differential likelihood of having non-missing and positive income. This is encouraging not only for interpreting the log income results, but also supports the notion that children of covered workers were not more likely to subsequently enter

[29]For the period 1937–1950 we only observe total Social Security earnings in nominal values. After 1950 we observe Social Security earnings separately by year, and convert them into 1972 US dollars before aggregating.

covered employment themselves. In terms of magnitude, a ten-year difference in father's minimum age of Social Security eligibility is associated with a roughly 100 percent increase in Social Security earnings between 1937 and 1960, and a roughly 50 percent increase in 1972 income measures. While these magnitudes are admittedly large, we emphasize that the OCG results should be viewed as suggestive and interpreted with caution given the small sample sizes involved, and the associated statistical imprecision surrounding the estimates.

The right panel of Figure 8 looks at three measures of occupational standing: the occupational income score associated with sons' first job, the job they held in 1962, and their job at the time of the survey in 1973.[30] The advantage of occupational standing measures is that they are less prone to transitory income shocks and are available even for workers whose earnings were not reported to the SSA or the IRS (though some individuals have missing occupations but non-missing income). The effect on the occupational income score of sons' first job is negative but not statistically significant. The coefficients for sons' occupational income scores in 1962 and 1973 are also negative, but slightly larger in magnitude and statistically significant at the 5 percent level. A ten-year difference in father's minimum age of Social Security eligibility is associated with a roughly 30 percent increase in sons' occupational income score during their prime to late adulthood.

6 Implications

6.1 Social Security conferred large financial gains on recipients' children.

Our results suggest that Social Security had substantial, long-lasting effects on the children of recipients. The most direct evidence of the financial gains from Social Security to recipients' children is the improvement in the quality of the neighborhoods in which they lived late in life. We estimate that a ten-year difference in the predicted earliest age of father's eligibility—roughly the difference in the actual earliest age of eligibility between 1935- and 1950-covered fathers from the 1877 birth cohort—is associated with the average son's ZIP at death having 2.3% higher average income and 2.8% higher median house prices. If recipients' children enjoyed the same 2.3% increase in income starting when their fathers were 65 years old, then their financial gain would have been equivalent to a $640 per year annuity or roughly $12,800 in expected present value, in 2000 dollars.[31,32]

[30] Occupational income scores are obtained by mapping OCG occupations into 1950 Census occupations and assigning to each occupation its corresponding median total income in the 1950 Census (Ruggles et al., 2021).

[31] These estimates come from discounting a gain of 2.3% of average family income in each calendar year starting in 1945 (when the 1880 cohort reaches age 65) to reflect time discounting (real interest rate of 3% per year) and mortality (based on the cohort life table for males in the 1900 cohort from the SSA).

[32] In general, people living in a location with higher income or higher house prices do not necessarily enjoy higher living standards. In simple spatial equilibrium models without individual heterogeneity, for example, everyone enjoys the same real living standard despite potentially large differences in nominal income and house prices. In reality,

What about daughters? The vast majority of the children of fathers affected by our instrument are married and so have two sets of parents, one set for each spouse. For a husband-and-wife couple, the effect of both of their fathers having earlier eligibility for Social Security is the sum of our estimates for sons and daughters. Given that the estimates for daughters tend to be small, we focus on the estimates for sons and treat them as applying to married couples as well.[33]

6.2 Key driver: Migration to better-matched labor markets

Our findings suggest that a key proximate factor underlying the financial gains from Social Security to recipients' children is that it increased their earnings. That recipients' children lived in ZIP codes with higher income and higher house prices is suggestive that they had higher earnings.[34] The estimates of individual-level earnings effects from the OCG, while imprecise, broadly corroborate this idea. What underlying factors led to these earnings gains? A variety of evidence suggests that the key driver was that Social Security led recipients' children to migrate to better-matched labor markets.[35]

One piece of evidence is that Social Security significantly increased migration to different labor markets, and such migration has been found to lead to large earnings gains in many contexts. For example, we estimate that a ten-year difference in the predicted earliest age of father's Social Security eligibility increased migration to a different commuting zone by roughly 2 percentage

such spatial equilibrium forces linking labor and housing markets likely operate mainly at the level of labor markets (e.g., CZs), not within labor markets (e.g., ZIP codes), so within-labor-market differences in income and house prices likely mainly reflect differences in real living standards across different types of households. As discussed in more detail below and consistent with spatial equilibrium forces toward the equalization of real living standards across labor markets, we find little association between Social Security and CZ characteristics, so the strong association between Social Security and ZIP characteristics comes from within-CZ sorting to better ZIP codes rather than across-CZ sorting to CZs with higher income or house prices. The strong association between Social Security and ZIP characteristics thus likely reflects differences in real living standards.

[33] Given that the vast majority of the children of fathers affected by our instrument are married, we interpret our estimates as showing that for husband-and-wife married couples, the Social Security of the husband's father is a more important driver of their migration and living standards than the Social Security of the wife's father. This is in keeping with the idea that sons and daughters-in-law play a different role in family old-age support arrangements than daughters and sons-in-law do, with the role of sons and daughters-in-law being more substitutable with Social Security benefits, perhaps because such couples tend to provide more resource-intensive support whereas daughters and sons-in-law couples provide more time- or care-intensive support. The small effects on daughters of their father's Social Security does not imply that unmarried daughters gain little from their father's Social Security. Unfortunately, it is not straightforward to estimate heterogeneity in the effects of Social Security by marital status because marriage might be affected by father's Social Security.

[34] In principle, this could also reflect, at least in part, higher non-labor income due to higher net transfers from their parents as a result of Social Security. In practice, as we soon discuss, the effects on ZIP code characteristics far exceed what would be expected from higher net transfers from their parents alone.

[35] Another potential driver is Social Security leading to a substitution among recipients' children from time spent supporting one's parents to time spent in market work. We return to this below. Note that income effects from recipients' children receiving larger net transfers from their parents as a result of Social Security are an opposing force that would tend to make Social Security reduce the earnings of recipients' children.

points (see Figure 7).[36] This migration appears to have been driven not by overall wage levels in a labor market but rather by comparative advantage based on the quality of the match between the household and the labor market. For example, Social Security did not affect the type of CZs recipients' children lived in as measured by CZ average income (see Appendix Figure A.13). This is what would be expected in standard spatial equilibrium models (Rosen, 1979; Roback, 1982), and it is consistent with recent evidence on the key role of comparative advantage in driving the gains from migration (e.g., Nakamura et al., 2021).[37]

If the financial gains from Social Security-induced migration were anywhere near as large as the gains from many other types of migration that have been studied in the literature, then the financial gains from Social Security-induced migration were large both in absolute terms and as a share of the overall financial gain from Social Security to recipients' children. For example, using very different approaches in very different settings, Kennan and Walker (2011) (estimated structural model for young men in the U.S.) and Nakamura et al. (2021) (quasi-experimental evidence from a volcanic eruption in Iceland) both find migration gains on the order of 100% of income. More generally, the broader literature on migration suggests that the financial gains from migration are often quite large.

As an illustrative benchmark, suppose that Social Security-induced migration to a new labor market conferred a gain of 50% of mean family income in each year from age 40 (father age around 65) to 64 and 25% thereafter. Under this benchmark, the Social Security-induced increase in migration would account for 39% (56%) of the overall financial gain to recipients' children if labor markets are defined as CZs (counties).

Of course, a gain of 50% of mean income from migrating to a new labor market, while smaller than prominent estimates in the literature, is still a large gain. The actual gain from this particular channel—greater net migration to a new labor market—might have been smaller. But this is not the only channel by which Social Security likely increased the gains to recipients' children from migration. Another channel is along the intensive margins of the timing and distance of migration. Social Security, in reducing the demand for family old-age support, may have caused migration among recipients' children to occur earlier in the life cycle (e.g., at the beginning of adulthood

[36] Note that commuting zones are based on commuting patterns in the 1990 Census, long after the prime working years of children in our sample. Given the improvements in transportation technology in the intervening decades, especially expansions of and improvements to road networks, the relevant labor markets for recipients' children were likely significantly smaller than CZs (see Pooley and Turnbull, 1999, on the expansion of labor markets during this time period, albeit in the U.K.). As such, in this time period migration to a new CZ likely understates the extent of migration to new labor markets.

[37] Social Security caused a large reduction in the probability of the joint outcome of staying in parents' 1930 CZ and living in a ZIP code in the bottom or second quintile of income or house prices (see Appendix Figure A.14). This is in keeping with the idea that with little Social Security, some recipients' children stay in their 1930 CZ despite being a poor match for that labor market, though it might also reflect that Social Security decreases net upstream transfers of money and time.

rather than after one's parents' old age) or to farther-away locations (e.g., by reducing the likelihood of returning periodically or permanently to provide old-age support). Given the high rates of migration during this time period—48% of recipients' children left the 1930 CZ and 63% left the 1930 county—even modest increases in the average return to such migration would confer large gains. For example, if Social Security increased the average return to migration away from the 1930 CZ (county) by 3% of mean family income from age 40–64 and 1.5% thereafter, such a channel alone would explain 56% (73%) of the overall financial gain to recipients' children.

A complementary approach to assessing the likely importance of migration in driving the financial gains from Social Security to recipients' children is to quantify the likely *non*-migration gains, e.g., from increasing the net financial transfers they receive from their parents. As we discuss in Appendix Section A.4, our calculations suggest that non-migration gains are unlikely to account for more than about one-fifth of the overall financial gains from Social Security to recipients' children. This suggests that migration gains likely account for at least four-fifths of the financial gains to recipients' children.[38]

Hence, consideration of the likely gains to recipients' children from (i) Social Security-induced changes in migration and (ii) non-migration factors (e.g., Social Security-induced changes in financial transfers between recipients and their children) both suggest that migration to better-matched labor markets accounted for most of the substantial financial gains from Social Security to recipients' children.

6.3 Incidence and efficiency of Social Security

Our findings suggest that Social Security conferred large financial gains on recipients' children in large part by leading them to migrate to better-matched labor markets. The size of these gains suggests important modifications to standard analyses of the incidence and efficiency effects of Social Security.

In terms of intergenerational incidence, our estimates suggest that the average gain to recipients' children from Social Security expansions exceeded the average gain to recipients themselves.[39] Hence, contrary to prevailing views, Social Security expansions do not appear to have

[38]These calculations account for two key ways in which Social Security might confer non-migration gains on recipients' children: an increase in the net downstream financial transfers they receive from their parents and a decrease in the time they spend supplying old-age support to their parents. As a benchmark, we assume that the resulting financial gains total 50% of the associated Social Security benefits. Together with an estimate of the "first-stage" effect of our instrument on the expected present value of Social Security benefits, this assumption implies non-migration gains of about one-fifth of the estimated overall per-child gain.

[39]Our benchmark estimate is that the net financial gain to recipients' children was twice the corresponding Social Security benefits, itself a likely upper bound on the net financial gain to recipients (given that Social Security benefits tend to decrease net upstream family transfers). Given such a large difference in net financial gains, it is unlikely that any differential consumption-smoothing benefits or non-pecuniary costs would overcome it. For example, the consumption-smoothing "markup" on Social Security benefits over their actuarial value would have to be over 100%

entailed a large redistribution from younger to older cohorts but rather largely a replacement of one type of old-age support (informal, mostly in-kind support within families) with another (formal, cash support through Social Security taxes and benefits). This replacement conferred large financial gains on recipients' children in large part by facilitating their migration to better-matched labor markets.

In terms of *intra*-generational incidence, our results imply that Social Security systematically redistributed from larger to smaller families. This can be viewed either as (ex post) redistribution across different "types" or (ex ante) insurance against risk in the costs of providing old-age support from behind the veil of ignorance. This arose both because larger families have more working-age taxpayers per retirement-age benefit recipient and because the relaxation of location constraints from displacing family old-age support tends to confer a bigger benefit per child to smaller families, since many of the costs of family old-age support are relatively fixed and since, other things equal, the lowest-cost provider of family old-age support has a higher cost on average in smaller families.

The patterns of both the inter- and intra-generational effects of Social Security are suggestive that the key driver was the displacement of in-kind family old-age support that required physical proximity, i.e., a reduction in upstream in-kind transfers from children to their elderly parents (e.g., shared housing, chores, and personal care), rather than an increase in net downstream financial transfers. In terms of intergenerational incidence, the financial gain from Social Security to recipients' children is, as discussed above, significantly larger than what plausible increases in net downstream financial transfers would produce but is consistent with significant earnings gains from migrating to better-matched labor markets in response to a reduction in demand for in-kind family old-age support. In terms of intra-generational incidence, the larger effects on sons than on daughters is also suggestive of displacement of family old-age support rather than an increase in net downstream financial transfers. Whereas an increase in net downstream financial transfers might be expected to confer a similar benefit on sons and daughters, a Social Security-induced reduction in upstream in-kind transfers from children to their elderly parents would likely benefit sons more than daughters given that sons appear to have been more likely to provide the type of old-age support more readily replaced by Social Security (e.g., financial support in the form of shared housing rather than personal care).[40,41]

for this factor to equalize the gains to recipients and their children. But such a markup is significantly higher even than that on unemployment insurance benefits (a leading estimate of which is 60%; see Hendren, 2017), and the markup on Social Security seems unlikely to be as high as that on UI (since unemployment is in many ways harder to self-insure than retirement: less predictable timing, earlier in the life cycle when liquidity constraints are more binding, etc.).

[40]The other key pattern of intra-generational incidence, the larger per-child effects for children in smaller families, is similarly consistent with Social Security reducing upstream in-kind transfers and increasing net downstream financial transfers and so does not obviously favor one over the other.

[41]The greater effects of fathers' Social Security on sons than on daughters does not imply that Social Security redistributed from women to men. Many women and men are part of couples, so to the extent that couples share resources, women and men both gain from Social Security transfers to either set of their parents. Our finding that

In terms of efficiency, our findings are suggestive that Social Security, by displacing family old-age support, reduced the effective migration costs of recipients' children and thereby led to a more efficient matching of recipients' children to labor markets. The associated efficiency gain is an important benefit of Social Security.[42] Under a benchmark estimate that migration accounted for four-fifths of the overall gain from Social Security to recipients' children, we estimate that a ten-year reduction in the predicted earliest age of eligibility produced a migration gain of about $47,000 of expected present value. This is over twice our benchmark estimate of the associated difference in the expected present value of Social Security benefits. In other words, under this benchmark the average dollar of Social Security benefits led to a migration-driven increase in earnings of more than $2.[43]

Even conservative estimates of the financial gain from Social Security-induced migration exceed plausible levels of other sources of efficiency gains from Social Security that have been emphasized in the literature. Consider the potential consumption-smoothing gain from correcting myopic failure to plan for retirement or from imperfections in private annuity markets. Although there is little direct evidence on the magnitude of this gain, it seems unlikely to be even half as large as the migration gains.[44] In other words, the financial gain from Social Security facilitating a more efficient matching of recipients' children to labor markets is a crucial determinant of its overall efficiency effect.

The private and fiscal externality gains from the induced effects on the migration and earnings of recipients' children may be a crucial determinant of the overall gain from old-age support

Social Security led sons and their families to move to different CZs is consistent with the idea that the type of old-age support provided by sons and their spouses (daughters-in-law) is more substitutable with the cash support that Social Security provides than is the support provided by daughters and their spouses (sons-in-law), which is consistent with daughters providing more time-intensive, care-based support and sons providing more money-intensive support. That we can separately identify the effects of a father's Social Security coverage on his sons versus daughters is an advantage of the fine-grained policy variation we exploit. Common types of more aggregate policy variation, by contrast, would not separately identify the effects of the Social Security coverage of one's father versus one's father-in-law since those would not vary independently of each other.

[42]See Appendix Section A.5 for a discussion of how Social Security may have facilitated the exploitation of gains from trade in old-age support that were otherwise foregone due to missing markets and transaction costs. It discusses how Social Security likely smoothed the financial component of old-age support across individuals and over time, in a way that would have been hard to achieve otherwise, and how such a smoothing in turn likely reduced the opportunity cost of migrating away from one's parents. In this way, Social Security may have facilitated a more efficient allocation of people across places. Evidence in keeping with this idea includes that Social Security appears to have significantly increased the average total family earnings of recipients and their children.

[43]To be clear, this net financial gain to recipients' children may overstate the associated efficiency gain if some of this gain came at the expense of reductions in net upstream support to recipients or net non-pecuniary costs from reductions in family contact.

[44]The consumption-smoothing "markup" on Social Security benefits over their actuarial value seems unlikely to match that on unemployment insurance, since, as discussed in footnote 39, unemployment is in many ways harder to self-insure than retirement. Hendren (2017) estimates a markup on unemployment insurance benefits of 60%. Such a markup on Social Security benefits would make the consumption-smoothing gain less than 30% of our estimate of its migration gain (60% versus over 200% of Social Security benefits).

programs more generally, perhaps especially long-term care programs. Whereas Social Security seems likely to have mainly displaced family old-age support of everyday living expenses, long-term care continues to be provided mostly by family members, including adult children (see, e.g., Barczyk and Kredler, 2018). Coe et al. (2015) find that children whose parents are exposed to higher subsidies for private long-term care insurance are less likely to co-reside or live nearby their parents and are more likely to work full-time. This is consistent with the idea that the overall efficiency effects of long-term care-related policies may depend crucially on any induced changes in the migration and earnings of the adult children of the direct recipients. This is an important topic for future research.

Of course, the extent to which our key results generalize to other types of government old-age support or other contexts depends on many factors. Two that seem likely to be especially important are the extent to which government old-age support displaces family old-age support and the returns to migration to better-matched labor markets. On the displacement of family old-age support, certain types of family old-age support are less prevalent in the U.S. today (e.g., shared housing), but other types remain substantial (e.g., informal care for elderly parents with chronic health conditions; see Barczyk and Kredler, 2018). Moreover, the expansion of markets for certain types of old-age support, such as formal home care, may make it easier for families to substitute away from family-provided care in response to expansions of government old-age support. These differences might suggest that whereas we find evidence that Social Security reduced family support mainly from sons to their elderly parents (likely often in the form of shared housing), expansions of government old-age support today, especially of in-kind long-term care, might reduce family support mainly from daughters and mainly in the form of informal care.[45] In terms of the gains from migration, Kaplan and Schulhofer-Wohl (2017) present evidence that the geographic specificity of returns to different occupations declined toward the end of the 20th century, which would tend to reduce the returns to migration. The increase in female labor supply in the latter half of the 20th century likely changed the pattern of potential household-level gains from migration by increasing the return to being in thick labor markets (e.g., Costa and Kahn, 2000), with an ambiguous effect on the average return to migration (e.g., if members of couples tend to be sufficiently dissimilar in terms of their optimal labor markets, an increase in two-earner couples could reduce the returns to migration).

[45] Other potentially relevant factors include norms about supporting one's parents in their old age, which may have weakened, and the amount of government old-age support, which has increased dramatically. But while these would likely reduce the *level* of family old-age support, it is not obvious how they would affect its *responsiveness* to changes in government old-age support.

7 Conclusion

Understanding the effects of Social Security on behavior and welfare requires understanding its effects not only on direct recipients but also on their children. We explore this question by combining a novel empirical approach, based on the fact that Social Security coverage was not universal at first but later expanded, with a newly-built dataset linking Social Security recipients to the long-run outcomes of their children. We find that sons whose parents were more likely to have received Social Security, or received it earlier, were more likely to have moved from their parents' 1930 location, had higher incomes, and settled in higher socioeconomic status ZIP codes. Consistent with both historical and modern survey evidence that daughters tend to provide care that is not typically purchased on the market, we find no evidence of effects among daughters.

The key driver of the financial gains for sons appears to be migration to better-matched labor markets. The resulting gains are large enough that for the average family, the total financial gains to recipients' children likely exceeded the gains to recipients themselves. A variety of indirect evidence suggests that these effects were driven by Social Security displacing old-age support from children to their parents that required physical proximity. Together, our findings suggest that the impact of government old-age support programs on the allocation of people across locations may be central to their overall welfare effects.

References

Abramitzky, Ran, Leah Boustan, Katherine Eriksson, James Feigenbaum, and Santiago Pérez, "Automated Linking of Historical Data," *Journal of Economic Literature*, 2021, *59* (3), 865–918.

_ , Roy Mill, and Santiago Pérez, "Linking individuals across historical sources: A fully automated approach," *Historical Methods: A Journal of Quantitative and Interdisciplinary History*, 2020, *53* (2).

Altmeyer, Arthur J., "The Desirability of Extending Social Security to Employees of Nonprofit Institutions," *Social Security Bulletin*, 1944, *7* (8).

_ , "Social Security for State and Local Government Employees," *Social Security Bulletin*, 1945, *8* (4), 4–7.

American Hospital Association, "For Flexibility: Changes in Pension Plan By-Laws," *Hospitals: the Journal of the American Hospital Association*, 1949, *23* (2), 61.

Ardington, Cally, Anne Case, and Victoria Hosegood, "Labor Supply Responses to Large Social Transfers: Longitudinal Evidence from South Africa," *American Economic Journal: Applied Economics*, 2009, *1* (1), 22–48.

Attanasio, Orazio P. and Agar Brugiavini, "Social Security and Households' Saving," *The Quarterly Journal of Economics*, 2003, *118* (3), 1075–1119.

Bailey, Martha J., Connor Cole, and Catherine Massey, "Simple Strategies for Improving Inference with Linked Data: A Case Study of the 1850–1930 IPUMS Linked Representative Historical Samples," *Historical Methods: A Journal of Quantitative and Interdisciplinary History*, 2020, *53* (2), 80–93.

_ , Peter Z. Lin, Shariq Mohammed, Paul Mohnen, Mengying Zhang, and Alexa Prettyman, "The Longitudinal, Intergenerational Family Electronic Micro-Database (LIFE-M)," Working paper 2021.

Barczyk, Daniel and Matthias Kredler, "Evaluating Long-Tterm-Care Policy Options, Taking the Family Seriously," *The Review of Economic Studies*, 2018, *85* (2), 766–809.

Barro, Robert J., "Are Government Bonds Net Wealth?," *Journal of Political Economy*, 1974, *82* (6), 1095–1117.

Bau, Natalie, "Can Policy Change Culture? Government Pension Plans and Traditional Kinship Practices," *American Economic Review*, 2021, *111* (6), 1880–1917.

Becker, Gary S., "A Theory of Social Interactions," *Journal of Political Economy*, 1974, *82* (6), 1063–1093.

Bernheim, B. Douglas and Kyle Bagwell, "Is Everything Neutral?," *Journal of Political Economy*, 1988, *96* (2), 308–338.

Blau, Peter M., Otis Dudley Duncan, David L. Featherman, and Robert M. Hauser, *Occupational Changes in a Generation, 1962 and 1973 (ICPSR 6162)*, Ann Arbor, MI: Inter-university Consortium for Political and Social Research [distributor], 1994. https://doi.org/10.3886/ICPSR06162.v1.

Bugbee, George, "Pensions Plan Is Now Open for Enrollment," *Hospitals: the Journal of the American Hospital Association*, 1946, *20* (9), 59–60.

Carter, Susan B., Scott Sigmund Gartner, Michael R. Haines, Alan L. Olmstead, Richard Sutch, Gavin Wright, Price V. Fishback, and Melissa A. Thomasson, *Historical Statistics of the United States Millennial Edition Online*, New York: Cambridge University Press, 2006.

Chen, Xi, Karen Eggleston, and Ang Sun, "The impact of social pensions on intergenerational relationships: Comparative evidence from China," *The Journal of the Economics of Ageing*, 2018, *12*, 225–235.

Chyn, Eric and Lawrence F. Katz, "Neighborhoods Matter: Assessing the Evidence for Place Effects," *Journal of Economic Perspectives*, 2021, *35* (4), 197–222.

Clark, Robert L., Lee A. Craig, and Jack W. Wilson, *A History of Public Sector Pensions in the United States*, University of Pennsylvania Press, 2003.

_ , _ , and John Sabelhaus, *State and Local Retirement Plans in the United States*, Edward Elgar, 2011.

Coe, Norma B., Gopi Shah Goda, and Courtney Harold Van Houtven, "Family Spillovers of Long-Term Care Insurance," NBER Working Paper 21483, National Bureau of Economic Research, Inc 2015.

Coile, Courtney C., "Economic Determinants of Workers' Retirement Decisions," *Journal of Economic Surveys*, 2015, *29* (4), 830–853.

Costa, Dora L., "Displacing the Family: Union Army Pensions and Elderly Living Arrangements," *Journal of Political Economy*, 1997, *105* (6), 1269–1292.

_ , *The Evolution of Retirement* National Bureau of Economic Research Books, University of Chicago Press, 1998.

_ , "A House of Her Own: Old Age Assistance and the Living Arrangements of Older Nonmarried Women," *Journal of Public Economics*, 1999, *72* (1), 39–59.

_ **and Matthew E. Kahn**, "Power Couples: Changes in the Locational Choice of the College Educated, 1940–1990," *The Quarterly Journal of Economics*, 2000, *115* (4), 1287–1315.

Edmonds, Eric V., Kristin Mammen, and Douglas L. Miller, "Rearranging the Family?: Income Support and Elderly Living Arrangements in a Low-Income Country," *Journal of Human Resources*, 2005, *40* (1), 186–207.

Eggleston, Karen, Ang Sun, and Zhaoguo Zhan, "The Impact of Rural Pensions in China on Labor Migration," *The World Bank Economic Review*, 2018, *32* (1), 64–84.

Engelhardt, Gary V., Jonathan Gruber, and Cynthia D. Perry, "Social Security and Elderly Living Arrangements Evidence from the Social Security Notch," *Journal of Human Resources*, 2005, *40* (2), 354–372.

Feigenbaum, James, "Automated Census Record Linking: A Machine Learning Approach," Working paper 2016.

Feldstein, Martin and Jeffrey B. Liebman, "Social Security," in A. J. Auerbach and M. Feldstein, eds., *Handbook of Public Economics*, Vol. 4 of *Handbook of Public Economics*, Elsevier, 2002, chapter 32, pp. 2245–2324.

Ferrie, Joseph P., "A New Sample of Males Linked from the Public Use Microdata Sample of the 1850 US Federal Census of Population to the 1860 US Federal Census Manuscript Schedules," *Historical Methods: A Journal of Quantitative and Interdisciplinary History*, 1996, *29* (4), 141–156.

Fetter, Daniel, Lee Lockwood, and Paul Mohnen, "The Intergenerational Incidence of Government Old-Age Support: Evidence from the Early Social Security Era," in Mohnen, Paul Hong-Lum, *Essays in Labor Economics*, Northwestern University PhD Dissertation 2018.

Folbre, Nancy, ed., *For Love or Money*, Russell Sage Foundation, 2012.

Friedman, Jerome, Trevor Hastie, and Robert Tibshirani, *The Elements of Statistical Learning: Data Mining, Inference, and Prediction*, 2nd ed., Springer, 2009.

Goldstein, Joshua and Casey Breen, "Berkeley Unified Numident Mortality Database: Public Administrative Records for Individual Level Mortality Research," Working paper 2020.

Gratton, Brian, "The Poverty of Impoverishment Theory: The Economic Well-Being of the Elderly, 1890–1950," *The Journal of Economic History*, 1996, *56* (1), 39–61.

Grigoryeva, Angelina, "Own Gender, Sibling's Gender, Parent's Gender: The Division of Elderly Parent Care among Adult Children," *American Sociological Review*, 2017, *82* (1), 116–146.

Gruber, Jonathan, *Public Finance and Public Policy*, 4th ed., Macmillan, 2013.

Hamoudi, Amar and Duncan Thomas, "Endogenous Coresidence and Program Incidence: South Africa's Old Age Pension," *Journal of Development Economics*, 2014, *109* (C), 30–37.

Hayhow, Edgar C., "Pensions and Life Insurance Plans for Hospital Employees," *Hospitals: the Journal of the American Hospital Association*, 1940, *14* (11), 73–77.

Helgertz, Jonas, Steven Ruggles, John Robert Warren, Catherine A. Fitch, Ronald Goeken, J. David Hacker, Matt A. Nelson, Joseph P. Price, Evan Roberts, and Matthew Sobek, *IPUMS Multigenerational Longitudinal Panel: Version 1.0 [dataset]*, Minneapolis, MN: IPUMS, 2020.

Hendren, Nathaniel, "Knowledge of future job loss and implications for unemployment insurance," *American Economic Review*, 2017, *107* (7), 1778–1823.

Horowitz, Amy, "Sons and Daughters as Caregivers to Older Parents: Differences in Role Performance and Consequences," *The Gerontologist*, 1985, *25* (6), 612–617.

Huang, Wei and Chuanchuan Zhang, "The Power of Social Pensions: Evidence from China's New Rural Pension Scheme," *American Economic Journal: Applied Economics*, 2021, *13* (2), 179–205.

Jensen, Robert T., "Do private transfers 'displace' the benefits of public transfers? Evidence from South Africa," *Journal of Public Economics*, 2003, *88* (1-2), 89–112.

Kaplan, Greg and Sam Schulhofer-Wohl, "Understanding the Long-Run Decline in Interstate Migration," *International Economic Review*, 2017, *58* (1), 57–94.

Kennan, John and James R. Walker, "The Effect of Expected Income on Individual Migration Decisions," *Econometrica*, 2011, *79* (1), 211–251.

Kestenbaum, Bert, "State and Local Government Employees Covered Under Social Security, 1977-81," *Social Security Bulletin*, 1982, *45* (12), 11–13.

Leibowitz, George J., "Old-Age and Survivors Insurance: Coverage Under the 1950 Amendments," *Social Security Bulletin*, 1950, *13* (12), 3–10, 21.

Leimer, Dean R., "Cohort-Sepcific Measures of Lifetime Net Social Security Transfers," Technical Report, Social Security Administration 1994. Office of Research and Statistics Working Paper Number 59.

Manson, Steven, Jonathan Schroeder, David Van Riper, Tracy Kugler, and Steven Ruggles, *IPUMS National Historical Geographic Information System: Version 15.0 [dataset]*, Minneapolis, MN: IPUMS, 2020.

Marquis, James E., "Old-Age and Survivors Insurance: Coverage Under the 1954 Amendments," *Social Security Bulletin*, 1955, *18* (1), 3–10.

McCamman, Dorothy, "State and Local Employees Covered by Government Retirement Systems," *Social Security Bulletin*, 1943, *6* (7), 31–41.

_ , "Higher Educational Institutions and Social Security," *Social Security Bulletin*, 1948, *11* (1), 5–11.

_ , "Retirement Protection for State and Local Employees: Ten Years of Growth," *Social Security Bulletin*, 1953, *16* (5), 3–10,24.

_ , "State and Local Government Employment Under OASI, October 1955," *Social Security Bulletin*, 1956, *19* (2), 15–17.

McGarry, Kathleen and Robert F. Schoeni, "Social Security, Economic Growth, and the Rise in Elderly Widows' Independence in the Twentieth Century," *Demography*, 2000, *37* (2), 221–236.

Moffitt, Robert A., "Trends in Social Security Wealth by Cohort," in "Economic Transfers in the United States" NBER Chapters, National Bureau of Economic Research, Inc, December 1984, pp. 327–358.

Mohammed, Shariq and Paul Mohnen, "Black Economic Progress in the Jim Crow South: Evidence from Rosenwald Schools," Working paper 2021.

Mukherjee, Anita, "Intergenerational Altruism and Retirement Transfers: Evidence from the Social Security Notch," *Journal of Human Resources*, forthcoming.

Murray, Merrill G. and Ilse M. Smith, "Higher Educational Institutions and the Social Security Act," *Social Security Bulletin*, 1940, *3* (12), 21–28.

Nakamura, Emi, Jósef Sigurdsson, and Jón Steinsson, "The Gift of Moving: Intergenerational Consequences of a Mobility Shock," *The Review of Economic Studies*, 2021.

National Education Association, *Statutory Provisions for Statewide Retirement Systems*, Washington, D.C.: National Education Association of the United States, 1946.

Nelson, William, "Employment Covered Under the Social Security Program, 1935-84," *Social Security Bulletin*, 1985, *48* (4), 33–39.

Parker, Florence E., "Experience Under State Old-Age Pension Acts in 1935," *Monthly Labor Review*, 1936, *34* (4), 811–837.

Pérez, Santiago, "The (South) American Dream: Mobility and Economic Outcomes of First- and Second-Generation Immigrants in Nineteenth-Century Argentina," *The Journal of Economic History*, 2017, *77* (4), 971–1006.

Pooley, Colin G. and Jean Turnbull, "The Journey to Work: A Century of Change," *Area*, 1999, *31* (3), 281–292.

Roback, Jennifer, "Wages, Rents, and the Quality of Life," *Journal of Political Economy*, 1982, *90* (6), 1257–1278.

Rosen, Sherwin, "Wage-Based Indexes of Urban Quality of Life," in "Current Issues in Urban Economics," Johns Hopkins Univ. Press, 1979, pp. 74–104.

Ruggles, Steven, "Linking Historical Censuses: A New Approach," *History and Computing*, 2006, *14* (1–2), 213–224.

_ , "The Decline of Intergenerational Coresidence in the United States, 1850 to 2000," *American Sociological Review*, 2007, *72* (6), 964–989.

_ **, Sarah Flood, Sophia Foster, Ronald Goeken, Jose Pacas, Megan Schouweiler, and Matthew Sobek**, *IPUMS USA: Version 11.0 [dataset]*, Minneapolis, MN: IPUMS, 2021.

Schieber, Sylvester J. and John B. Shoven, *The Real Deal: The History and Future of Social Security*, New Haven: Yale University Press, 1999.

Schorr, Alvin Louis, *Filial responsibility in the modern American family: an evaluation of current practice of filial responsibility in the United States and the relationship to it of social security programs*, Washington, D.C.: U.S. Department of Health, Education, and Welfare, Social Security Administration, Division of Program Research., 1960.

Schottland, Charles I., "Social Security Amendments of 1956: A Summary and Legislative History," *Social Security Bulletin*, 1956, *19* (9), 3–31.

Shanas, Ethel, *Family Relationships of Older People: Living Arrangements, Health Status, and Family Ties of Those Aged 65 and Over, As Reportd by the Aged, the Persons to Whom They Would Turn in a Health Crisis, and the General Public*, New York: Health Information Foundation, 1961.

Sharp, Harry and Morris Axelrod, "Mutual Aid Among Relatives in an Urban Population," in Ronald et al. Freedman, ed., *Principles of Sociology: A Text with Readings*, New York: Henry Holt and Company, 1956, pp. 433–439.

Social Security Administration, "Social Security Bulletin, Annual Statistical Supplement 1977-79," 1980. U.S. Government Printing Office.

_ , *Current Population Survey, 1973, and Social Security Records: Exact Match Data (ICPSR 7616)*, Ann Arbor, MI: Inter-university Consortium for Political and Social Research [distributor], 2005. https://doi.org/10.3886/ICPSR07616.v1.

_ , *Social Security Longitudinal Earnings Public Use File, 1937-1975 (ICPSR 7617)*, Ann Arbor, MI: Inter-university Consortium for Political and Social Research [distributor], 2008. https://doi.org/10.3886/ICPSR07617.v1.

SSA Office of the Chief Actuary, "OASDI Trustees Report," 2013.

Steiner, Peter Otto and Robert Dorfman, *The Economic Status of the Aged*, Berkeley: University of California Press, 1957.

Streib, Gordon F., "Family Patterns in Retirement," *Journal of Social Issues*, 1958, *14* (2), 46–60.

Subcommittee on Social Security, "Background Material on Social Security Coverage of Governmental Employees and Employees of Nonprofit Organizations," in "Coverage and Termination of Coverage of Government and Nonprofit Organization Employees under the Social Security System: Public Hearings before the Subcommittee on Social Security of the Committee on Ways and Means, U.S. House of Representatives, Ninety-Fourth Congress, Second Session," U.S. Government Printing Office, 1976.

Sussman, Marvin B., "The Help Pattern in the Middle Class Family," *American Sociological Review*, 1953, *18* (1), 22–28.

_ , "Relationships of Adult Children with their Parents in the United States," in Ethel Shanas and Gordon F. Streib, eds., *Social Structure and the Family: Generational Relations*, Englewood Cliffs, N.J.: Prentice-Hall, Inc., 1965, pp. 62–92.

Tolbert, Charles M. and Molly Sizer, "US Commuting Zones and Labor Market Areas: A 1990 Update," Staff Paper no. AGES–9614, U.S. Department of Agriculture, Economic Research Service 1996.

United States Bureau of the Census, *Retirement Systems for State and Local Government Employees: 1941*, U.S. Department of Commerce, Bureau of the Census, 1943.

_ , *Retirement Coverage of State and Local Government Employees*, Washington, D.C.: U.S. Department of Commerce, Bureau of the Census, 1953.

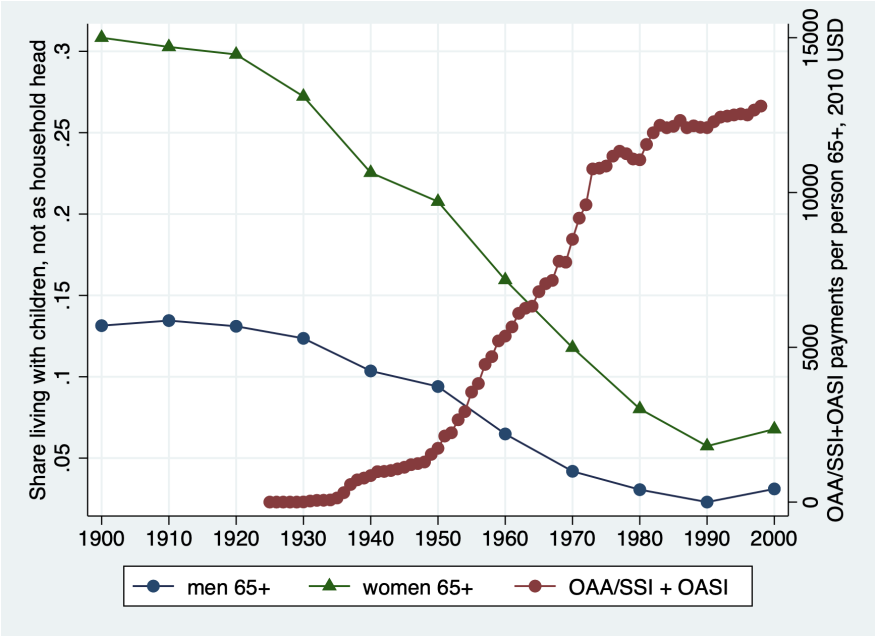
_ , *Current Population Survey: Annual Demographic File, 1973 (ICPSR 7564)*, Ann Arbor, MI: Inter-university Consortium for Political and Social Research [distributor], 1992. https://doi.org/10.3886/ICPSR07564.v1.

_ , "National Intercensal Tables, 1900–1990: Population by age, sex, and race," https://www.census.gov/data/tables/time-series/demo/popest/pre-1980-national.html 2004. Accessed: 2020-08-10.

Zimran, Ariell, "Sample-Selection Bias and Height Trends in the Nineteenth-Century United States," *The Journal of Economic History*, 2019, *79* (1), 99–138.

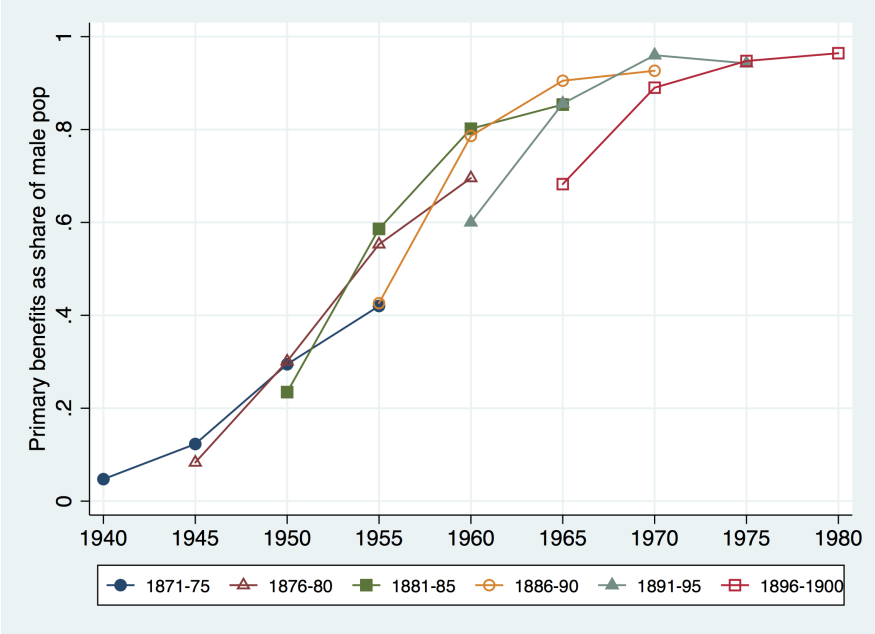
Figures and Tables

Figure 1: Intergenerational co-residence over the 20th century



Notes: This graph plots the share of men and women aged 65 or older living with one of their children or children-in-law, not as a household head or spouse of the household head. Also displayed are total payments under Old Age Assistance (OAA) and Old Age and Survivors Insurance (OASI) divided by the 65+ population (in 2010 US dollars). OAA payments data come from Parker (1936) for 1925 to 1935 and Series Bf621 of Carter et al. (2006) for 1936 onwards. OASI payments data come from Series BF396 of Carter et al. (2006). Data on co-residence come from Ruggles et al. (2021).

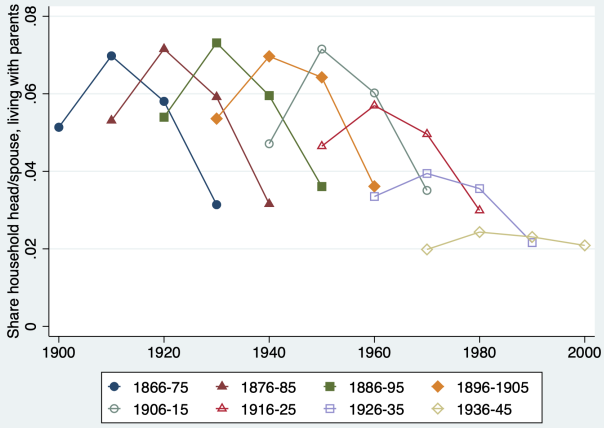
Figure 2: Receipt of retirement benefits by birth cohort and year



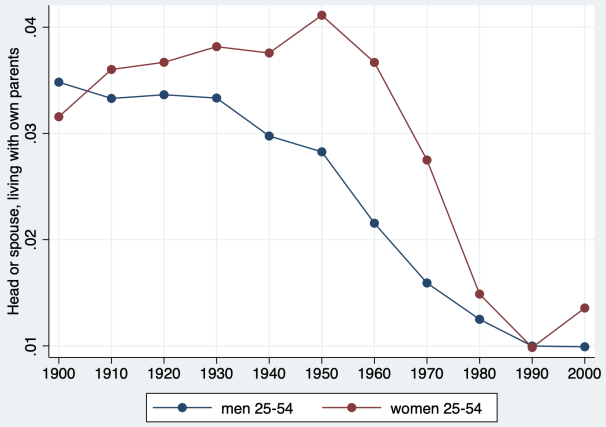
Notes: Figure shows share of men of each birth cohort receiving retirement benefits in the specified year, at ages 65–69, 70–74, 75–79, and 80–84. Data on payments are from the *Social Security Bulletin*, data on male population by age and year are from United States Bureau of the Census (2004).

Figure 3: Co-residence from children's perspective

(a) All children, by birth cohort



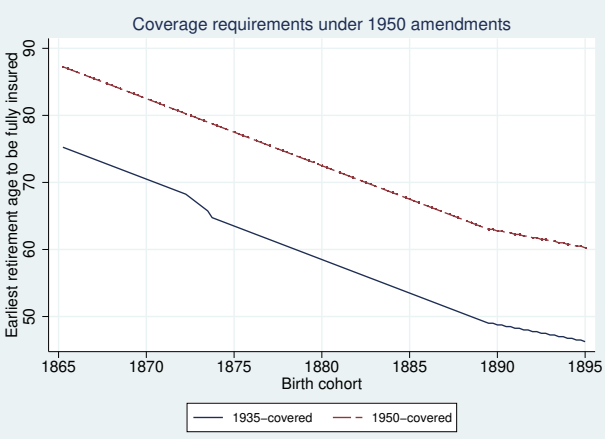
(b) Co-residence rates with own parents, by sex



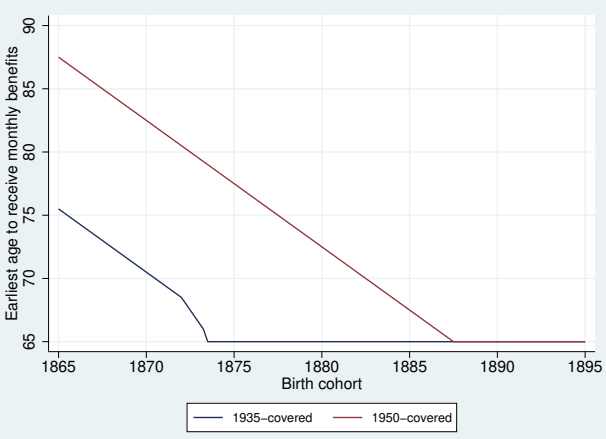
Notes: Panel (a) shows share of US-born men and women of each birth cohort who are household head or spouse of the head, and living with a parent or parent-in-law. Each cohort is shown at ages 25–34, 35–44, 45–54, and 55–64. Panel (b) shows share of US-born men 25–54 and US-born women 25–54 who were household head or spouse of the head and lived with their own parents in the household in each year. Data come from Ruggles et al. (2021).

Figure 4: Earliest age when permanently insured and earliest age for benefits

(a) Earliest age to be permanently insured

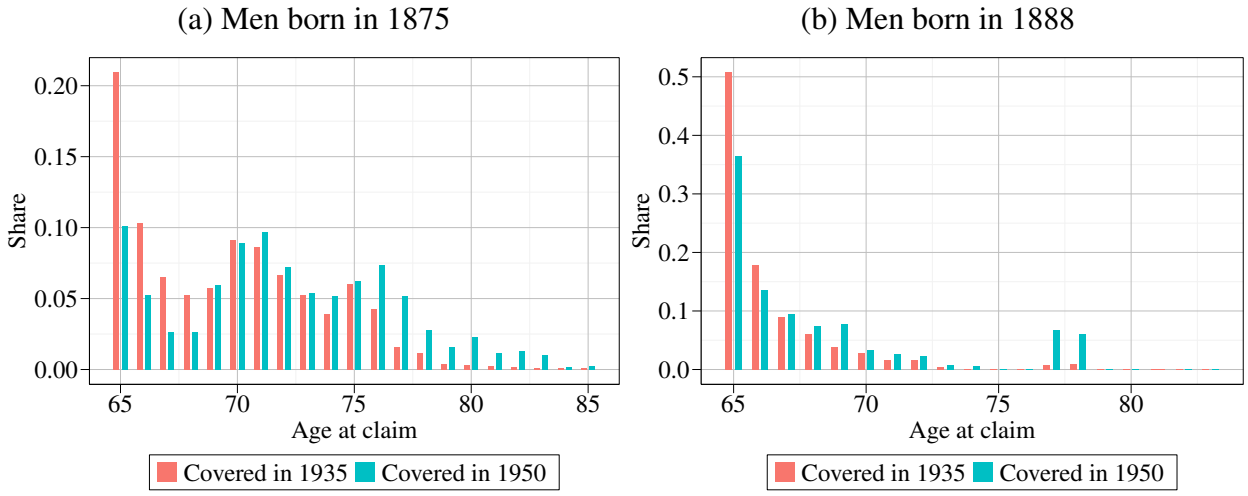


(b) Earliest age to receive monthly benefits



Notes: Minimum ages based on Social Security Act of 1935, 1939 Amendments, and 1950 Amendments, assuming that a worker in 1950-covered employment has no employment that was covered in 1935. The minimum ages shown are those under the 1950 Amendments.

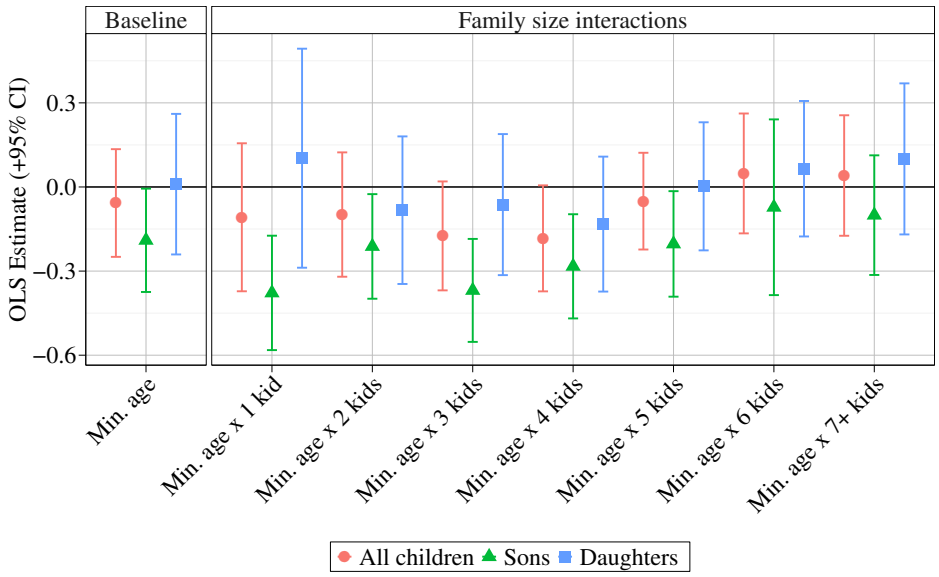
Figure 5: Age at claim distribution by Social Security coverage, as implied by 1930 employment



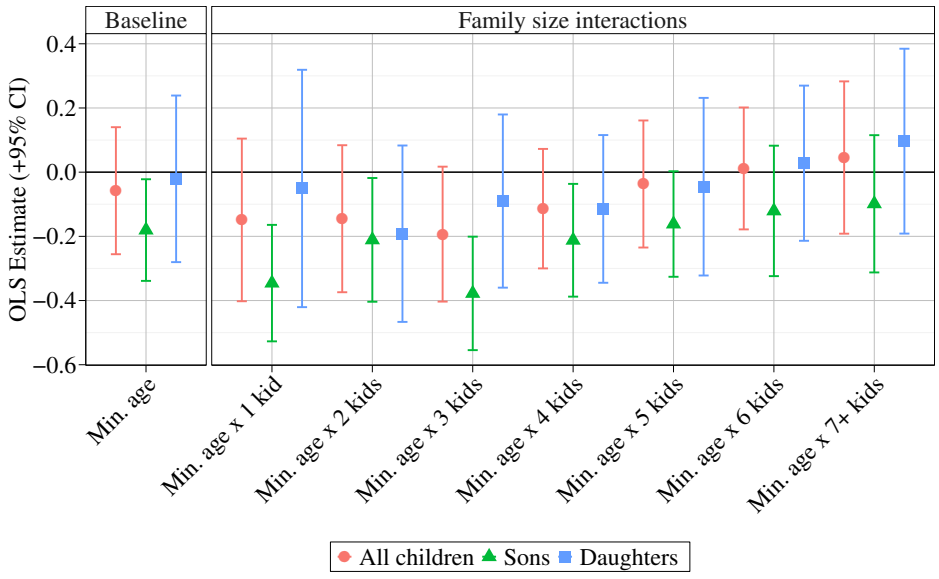
Notes: These graphs are based on the Numident claims-census linked sample, and plot the age at claim distribution for men whose 1930 employment was covered by Social Security in 1935 vs. 1950 (excluding the self-employed), separately for men born in 1875 and men born in 1888. Observations are weighted using IPWs.

Figure 6: The impact of father's later Social Security coverage on children's death ZIP rank

(a) Death ZIP rank in terms of 2001 mean AGI



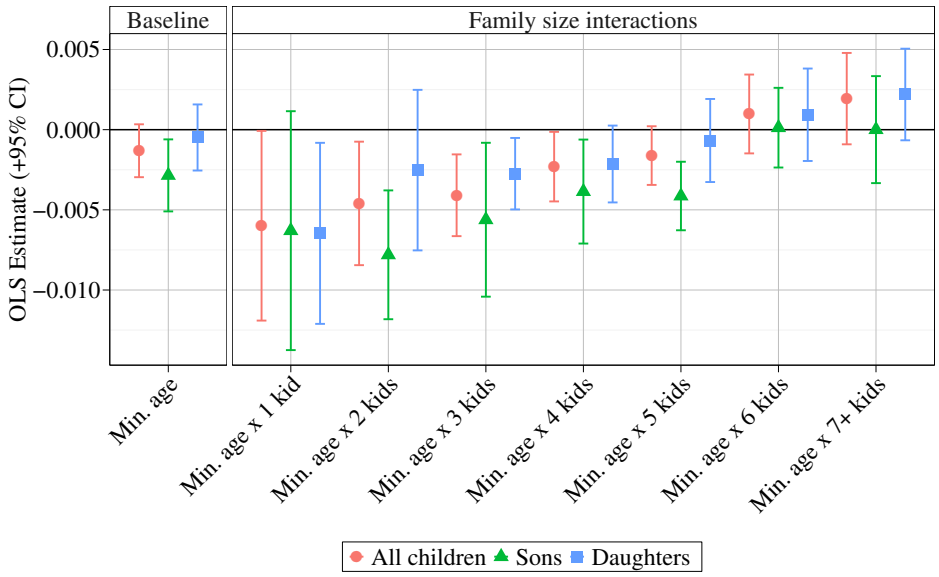
(b) Death ZIP rank in terms of 2000 median house values



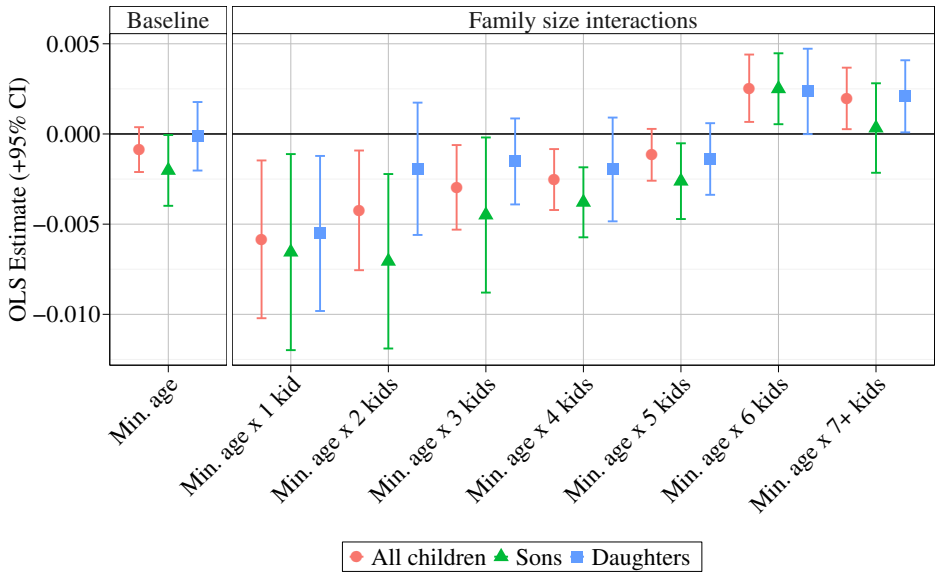
Notes: Figures display coefficients on father's minimum age of Social Security eligibility from equation (2), where the outcome is the average percentile rank of his children's ZIP code at death in terms of mean AGI or median house values (panel title). Baseline estimates shown in the left panel of each graph; estimates interacting minimum age with fixed effects for number of children shown in the right panel. Colors/shapes indicate separate estimates for all children, sons only, or daughters only (legend). All regressions include fixed effects for father's occupation × year of birth, Social Security coverage category, race, and state/country of birth. Observations weighted using IPWs. 95% confidence intervals based on standard errors clustered at the level of father's industry of employment.

Figure 7: The impact of father's later Social Security coverage on children's long-run migration

(a) Died outside parents' 1930 county of residence

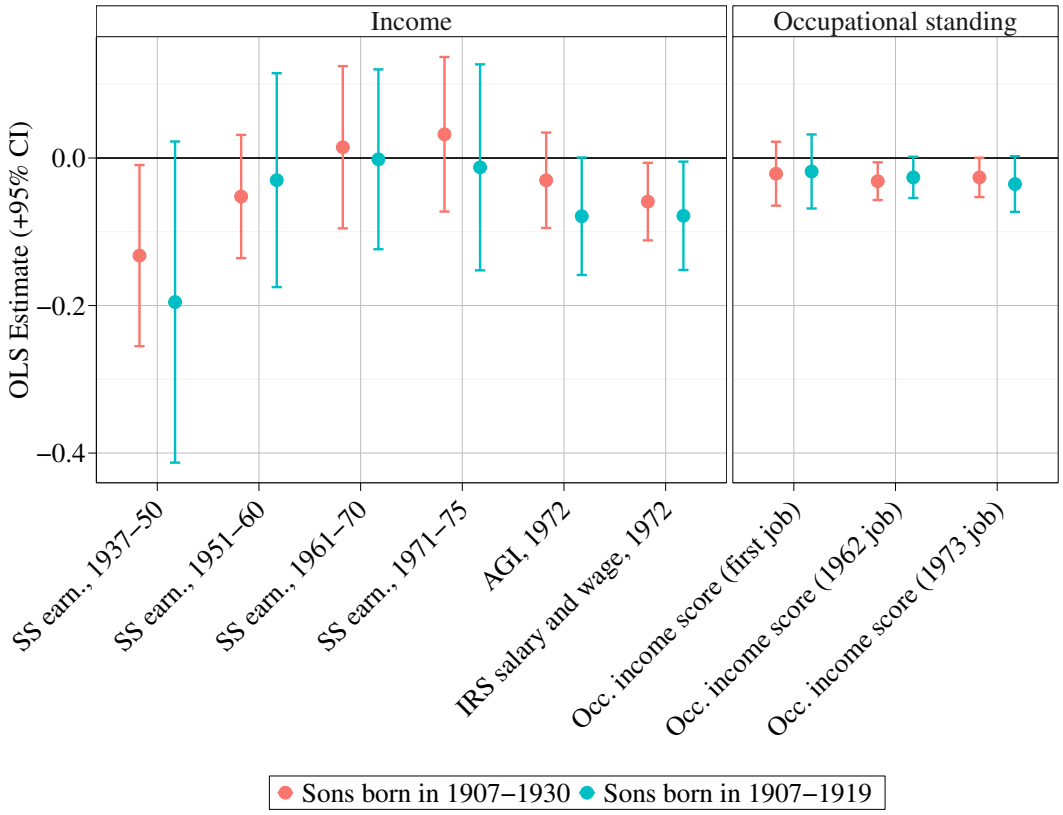


(b) Died outside parents' 1930 CZ of residence



Notes: Figures display coefficients on father's minimum age of Social Security eligibility from equation (2), where the outcome is the share of his children who died outside his 1930 county or CZ of residence (panel title). Baseline estimates shown in the left panel of each graph; estimates interacting minimum age with fixed effects for number of children shown in the right panel. Colors/shapes indicate separate estimates for all children, sons only, or daughters only (legend). All regressions include fixed effects for father's occupation × year of birth, Social Security coverage category, race, and state/country of birth. Observations weighted using IPWs. 95% confidence intervals based on standard errors clustered at the level of father's industry of employment.

Figure 8: The impact of father's later Social Security coverage on sons' income and occupational standing (in logs)



Notes: Figure displays coefficients on father's minimum age of Social Security eligibility from equation (2), where the outcome is indicated on the *x*-axis (all outcomes in logs). Colors/shapes indicate separate estimates for sons born in 1907–1930 and sons born in 1907–1919 (legend). All regressions include fixed effects for father's occupation × year of birth, Social Security coverage category, and son's year of birth, race, and state/country of birth. Observations weighted using OCG sampling weights. 95% confidence intervals based on standard errors clustered at the level of father's industry of employment.

Table 1: The impact of later Social Security coverage on subsequent claiming behavior

	Dependent variable:	
	1[linked to claims data]	Age at claim
	(1)	(2)
Min. age of Social Security eligibility	-0.0033***	0.1163***
	(0.0004)	(0.0237)
Sample	Population	Claims data
Mean of dep. var.	0.1	67.79
R^2	0.0234	0.1682
N	3,796,856	364,278

Notes: Unit of observation is a man in 1930 born in 1875–1888 and covered by Social Security in 1935 or 1950 based on his employment information, excluding self-employed workers. All regressions include fixed effects for occupation × year of birth, Social Security coverage category, race, and state/country of birth. Observations weighted using IPWs in column (2). Robust standard errors in parentheses, clustered at the industry level. *** 1%, ** 5%, * 10% significance.

Table 2: Selection of 1930 households and children into our linked sample

	Dependent variable: Any/number of children linked					
	All children		Sons		Daughters	
	Any	Number	Any	Number	Any	Number
	(1)	(2)	(3)	(4)	(5)	(6)
Min. age of Social Security eligibility	-0.0006	-0.0014	-0.0006	-0.0012	-0.0002	-0.0002
	(0.0006)	(0.0013)	(0.0006)	(0.0009)	(0.0003)	(0.0005)
Mean of dep. var.	0.216	0.362	0.136	0.178	0.137	0.184
R^2	0.0188	0.0178	0.0126	0.0118	0.0126	0.0124
N	2,971,208	2,971,208	2,971,208	2,971,208	2,971,208	2,971,208

Notes: Unit of observation is a 1930 household headed by a man born in 1875–1888 and covered by Social Security in 1935 or 1950 based on his employment information, excluding self-employed workers. The dependent variable is either an indicator equal to one if at least one child of a specific sex was linked to the 1930 household, or a discrete variable capturing the number of children of a specific sex linked to the 1930 household. All regressions include fixed effects for head of household occupation × year of birth, Social Security coverage category, race, and state/country of birth. Robust standard errors in parentheses, clustered at the head of household industry level. *** 1%, ** 5%, * 10% significance.

Table 3: Placebo tests: 1930 Census household characteristics

	Dependent variable: Household characteristic in 1930								
	Homeowner	House value (asinh)	Own a radio	Live in urban area	Live on farm	Number of children	Dad literate	Mom literate	Mom in labor force
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel (a): All families in population									
Min. age of Social Security eligibility	-0.0005	-0.0050	0.0003	-0.0003	-0.0002	-0.0038	-0.0001	-0.0000	0.0003
	(0.0012)	(0.0112)	(0.0015)	(0.0009)	(0.0002)	(0.0042)	(0.0006)	(0.0008)	(0.0009)
Mean of dep. var.	0.48	3.96	0.44	0.71	0.06	2.22	0.93	0.93	0.11
R^2	0.0640	0.0884	0.2154	0.2562	0.2332	0.1078	0.1495	0.1908	0.0625
N	2,964,018	2,899,011	2,971,208	2,971,208	2,971,208	2,971,208	2,971,208	2,708,885	2,709,477
Panel (b): All families in linked sample									
Min. age of Social Security eligibility	-0.0009	-0.0011	-0.0003	0.0003	-0.0005***	-0.0007	0.0001	0.0003	0.0003
	(0.0011)	(0.0110)	(0.0012)	(0.0006)	(0.0002)	(0.0031)	(0.0005)	(0.0006)	(0.0013)
Mean of dep. var.	0.48	4	0.44	0.7	0.06	3.95	0.93	0.93	0.09
R^2	0.0752	0.1045	0.2146	0.2602	0.2194	0.1292	0.1426	0.1822	0.0620
N	650,247	638,367	650,848	650,848	650,848	650,848	650,848	637,236	637,328

Notes: Unit of observation is a 1930 household headed by a man born in 1875–1888 and covered by Social Security in 1935 or 1950 based on his 1930 employment information, excluding self-employed workers. All regressions include fixed effects for head of household occupation × year of birth, Social Security coverage category, race, and state/country of birth. Observations weighted using IPWs in panel (b). Robust standard errors in parentheses, clustered at the head of household industry level. *** 1%, ** 5%, * 10% significance.

A Appendix

A.1 Details on linkages

A.1.1 Linking families in SS-5 records to households in the 1920 and 1930 Censuses

The first step involves reconstituting families in the SS-5 data by linking siblings together primarily based on parent names.[46] The sibling linkages are then used to create a crosswalk between SSNs and family identifiers. In practice, we borrow this crosswalk from Mohammed and Mohnen (2021). The sibling linkages were made using the LIFE-M supervised ML approach described in this paper, except that linking siblings together is an instance of one-to-many rather than one-to-one linking (i.e., individuals can have multiple siblings in the SS-5 data). As a result, trainers were allowed to select multiple links, and a random forest model was used to model their decisions. In particular, each primary-potential sibling pair was treated as a separate linking decision, even though trainers were technically allowed to select multiple siblings simultaneously. Further details can be found in Section B.2 from the online appendix of Mohammed and Mohnen (2021). The crosswalk contains 5.9 million families with at least two siblings, starting from 40 million unique SSNs. Alternatively, 35 percent of SSNs are linked to at least one other SSN in the SS-5 data.

The next step involves linking reconstituted families in SS-5 records to households in the 1930 Census. Although we are mainly interested in families in which the father was born in 1875–1888, since SS-5 records lack parental year of birth information we link all families with at least one child aged under 25 at the time of the census, and later restrict the sample based on age information from the census. Note that families can be a single SS-5 record if no sibling was found in the previous step. To generate the set of candidate census households for each family, we block on father last name initial and the *union* of children's state/continent of birth.[47] Blocking on the union of children's place of birth implies that the only requirement we impose is that at least one of the potential children's place of birth matches one of the primary children's place of birth. Within each block, we select the top 20 potential households based on a Jaro-Winkler string similarity score between primary and potential parent names.

We generate training data for a random sample of 500 primary families using the process described in Section 3. In particular, we display the following information to trainers: father full name, mother first and middle name, and children's first and middle name, age in 1930, and state/country of birth (we also display the modal race among children). We only display primary children that would have been alive at the time of the census, and mask potential children in the census if there is little chance that they match any of the primary children (based on age and place of birth considerations). We limit the number of children displayed to a maximum of eight to avoid overcrowding the screen during the training process (the eight youngest ones).

We are able to manually link 62 percent of families to a household in the 1930 Census in the training data. To train the two-stage model from Bailey et al. (2021) to make similar linking decisions, we first split the training data into a "training set" and a "test set" using a 70/30 percent split. The training set is used as input data to train the algorithm, while the test set can be used to assess the out-of-sample performance of the model (which typically lines up quite well with

[46] Race, place of birth, and year of birth information is also used in the linking decision, as siblings tend to be born in the same state/country a few years apart (and have matching race).

[47] Blocking on a subset of characteristics is common practice in the historical linking literature, and is necessary for computational reasons as string similarity scores are expensive to compute.

the in-sample performance). The set of features used to train the model is shown in Appendix Table A.8, and probabilities are generated using 10-fold cross-validation as described in Section 3. Panel (a) in Appendix Figure A.15 shows the resulting precision-recall curve. The model is able to achieve a recall rate of 85 percent at 97 percent precision. In other words, fixing the error rate of model-based links relative to human trainers at 3 percent, the model is able to replicate 85 percent of the links made by human trainers.

Using the trained model and cross-validated probability threshold, we scale-up linking to the full sample of SS-5 families. After doing so, it is not uncommon for the same census household to be linked to two (or more) different families in the SS-5 data. While some of these cases are true conflicts, a significant share of them are not. The reason is that we were not able to make all possible sibling linkages in the SS-5 data since parent names are not always consistently reported across siblings' SS-5 records. Although we might not be able to link two SS-5 records together solely based on parent name information, we still might be able to (correctly) link them to the same census household since the latter additionally exploits information on children. As a result, we end up with apparent conflicts that simply reflect the incompleteness of our family identifiers.

Rather than throwing away all these conflicts, we manually examined a random sample of them (200), and either chose one link over the others, concluded all links were correct, or chose none of them if we were not confident. We then used this data to assess the performance of different conflict resolution rules in terms of their ability to retain all correct links and drop all incorrect links.[48] The rule we ultimately settled on is able to retain 80 percent of correct links and drop 75 percent of incorrect links in the random sample of cases we manually inspected. The final number of SS-5 families linked to a household in the 1930 Census at 97 percent precision after resolving these conflicts is 6.9 million. Using an analogous procedure and recycling the training data and model, we are able to link 4.1 million SS-5 families to a household in the 1920 Census.

A.1.2 Linking households in the 1930 Census to households in the 1920 Census

The process of linking households in the 1930 Census to households in the 1920 Census is very similar to the one used to link SS-5 families to households in the 1930 Census. One slight difference is that the censuses additionally contain information on parents' age and place of birth. Therefore, when generating sets of potential candidates, we block on fathers' state/continent of birth and a +/-5-year window around fathers' age in 1920 (in addition to father last name initials). We also display father age in 1920, mother age in 1920, and mother place of birth information during the training process. We are able to manually link 53 percent of households headed by a man born in 1860–1905 in the 1930 Census to a household in the 1920 Census in our training data (500 cases). The model recall rate at 97 percent precision is 90 percent (see panel (b) in Appendix Figure A.15). Note that after scaling up the model it is possible for a household in the 1920 Census to be linked to multiple households in the 1930 Census. But in contrast to when we link SS-5 families to census households, these are true conflicts that primarily reflect name commonality. In some cases, we selected one link over the others based on perceived quality (as measured by the match probability). In the remaining cases, we dropped all conflicting links. The final number of households headed by a man born in 1860–1905 in the 1930 Census linked to the 1920 Census at 97 percent precision is 7.6 million.

[48] The two extreme rules are to keep all conflicts or drop all conflicts, which exactly achieve one objective but not the other. The optimal rule lies somewhere in between.

A.1.3 Linking individuals in claims records to individuals in the 1930 Census

Claims records for men born between 1865 and 1899 are linked to the 1930 Census using individual-level information. Given that our objective is to explore claiming behavior for regular retirement benefits, we focus on the subset of life claims.[49] We separately link claims records with non-missing place of birth information (40 percent of records) and claims records with missing place of birth information (60 percent of records). Claims records with non-missing place of birth information are linked using name, year of birth, and place of birth information. To generate sets of potential candidates, we block on state/continent of birth, last name initial, and a +/-3-year window around age in 1930. We then compute a string similarity score that accounts for the possibility of inverted first and middle names, and select the top 20 candidates. We manually linked a random sample of 1,000 cases, displaying name, age, and country of birth information (when applicable).[50] We also display the state associated with the first 3 digits of a primary's SSN and the state of residence of potential candidates in the 1930 Census.[51] Since 1930 is relatively close in time to the period in which these cohorts would have applied for an SSN, this optional piece of information is useful when choosing between multiple plausible candidates.

We are able to link 53 percent of claims records with non-missing place of birth information to an individual in the 1930 Census, and the model achieves a recall rate of 67 percent at 97 percent precision (see panel (c) in Appendix Figure A.15). We use an analogous procedure to link claims records with missing place of birth information, except that we cannot block on state/continent of birth, and instead block on the first two last name initials for computational reasons (in addition to the +/-3-year age window). This results in larger blocks and more ambiguity in the set of potential candidates. Consequently, we are only able to manually link 34 percent of claims records with missing place of birth information to an individual in the 1930 Census. The model recall rate at 97 percent precision is also slightly lower at 65 percent (see panel (d) in Appendix Figure A.15). After scaling up linking to the full set of claims records and resolving conflicts in which the same census individual was linked to two different claims records, we end up with a sample size of 1.3 million.

A.1.4 Linking OCG records to other survey/administrative records

Our goal here is to link four sets of records: (1) the 1973 wave of the Occupational Changes in a Generation survey (OCG), (2) the Social Security Records: Exact Match Data file (EMD), (3) the 1937–1975 Social Security Longitudinal Earnings File (SSLEF), and (4) the 1973 Current Population Survey (CPS). Our starting point is the OCG survey as it contains information on fathers' employment and year of birth. The other three datasets contain additional outcomes for sons. The EMD file contains tax record information from the Internal Revenue Service for the 1972 fiscal year, the SSLEF contains Social Security earnings records covering the period 1937–1975, and the CPS contains more granular 1973 place of residence information than in the OCG.

Unfortunately, there is no unique identifier to merge these datasets together. Instead, we proceed in three steps. First, we merge the EMD file with the SSLEF based on row order. More

[49] Life claims capture retirement and disability benefits. Death claims, which we exclude, capture survivor and death benefits. Around 77 percent of records in the claims data are life claims.

[50] For individuals born outside the U.S., only continent of birth matches, so an exact match on country of birth has additional informational content.

[51] Until 1972, the first 3 digits of an SSN was tied to the field office in which the application was filed.

specifically, after restricting the EMD file to individuals that should be found in the SSLEF, we end up with the same number of observations. Moreover, the observations seem to have been ordered in the same way in the raw data, as evidenced by the fact that demographic variables that appear in both datasets (sex, race, month of birth, and year of birth) match for over 99% of observations. Second, while the EMD file and the CPS file have the same number of rows, they are not ordered in the same way. Instead, we first impose the OCG sample restriction on both datasets (civilian males aged 20–65), and identify 24 variables that appear in both datasets and for which aggregate distributions line up exactly. These 24 variables uniquely identify over 99.9% of observations in both datasets. We therefore restrict the samples to uniquely-identified observations and merge them together. Similarly, we merge the EMD file and OCG file based on 34 variables that are contained in both datasets. The final merged OCG-EMD-SSLEF-CPS file contains 34,101 observations, or 99.9% of the original 34,142 observations in the OCG dataset.

A.2 Classification of employment into OASI coverage categories

A.2.1 Identifying Social Security coverage categories

Our assignment of Social Security coverage year to an individual is based on his reported occupation, industry, and "class of worker" (which includes information on whether a worker was self-employed, a wage or salary earner, or an unpaid family worker). Although our analysis uses only some of the Social Security coverage categories, we classify all combinations of occupation, industry, and class of worker and describe our approach here. We do not offer an exhaustive description of all categories, but cover the largest ones. Our description and classification is based primarily on Nelson (1985), supplemented with information from Marquis (1955) and Social Security Administration (1980), as well as the Social Security Act and subsequent amendments. Additional useful references include Leibowitz (1950) on the 1950 Amendments, Schottland (1956) on the 1956 Amendments, and McCamman (1956) on state and local government workers under the 1954 Amendments.

The original Social Security Act in 1935 included most wage and salary workers in its definition of covered employment. It excluded self-employment, as well as certain types of wage and salary work, the largest of which were agricultural wage work; domestic work in private homes; and work for nonprofit organizations, state and local government, the federal government, and railroads.

Since our analysis compares only wage and salary workers, we first describe the extension of Social Security coverage to the categories of wage and salary workers that were excluded from coverage in the 1935 Act. Regular agricultural wage work and domestic work were both covered under the 1950 Amendments, and the requirements for regularity were relaxed under the 1954 and 1956 Amendments. Employees of nonprofit organizations had been excluded from coverage under the 1935 Act because of the tax-exempt status of nonprofits (McCamman, 1948), but (with the exception of ministers) could be covered on an elective basis starting with the 1950 Amendments. The decision to elect coverage was made by both employers and employees: if the organization waived its exemption from Social Security taxes and employees voted to be covered by Social Security, all new employees and existing employees who voted for coverage would be covered, while existing employees who voted against coverage would not be covered. Although not all organizations elected coverage immediately upon its becoming available, take-up was reasonably rapid: as of 1951, 56% of employees of nonprofit organizations were covered under Social Security, rising

to 76% by 1960 (Subcommittee on Social Security, 1976, p. 27). State and local government employees were initially excluded from coverage because of perceived constitutional limitations on levying taxes on state and local governments (see, e.g. Clark et al., 2011), but beginning with the 1950 Amendments, states could voluntarily enter into agreements for Social Security to cover employees who were not covered under an existing retirement system. The 1954 Amendments extended coverage to state and local employees already under a retirement system, with the exception of firefighters and police officers, if the state elected to make coverage available and employees voted in favor of coverage. Amendments in 1956 further liberalized the provisions under which state and local government employees could be covered by Social Security.[52] About 11 percent of all state and local government employees were covered by Social Security at the beginning of 1951 (the first year in which they could be), and by 1958 the majority of state and local government employees (54 percent) were covered by Social Security (Kestenbaum, 1982).

Other significant types of work excluded from coverage under the 1935 Act, but that do not play a role in our analysis, are self-employment and work for railroads or the federal government. The 1950 Amendments covered most nonfarm self-employment, but excluded self-employed farmers and farm managers, as well as certain professional groups (accountants, architects, engineers, funeral directors, lawyers, medical professionals, and veterinarians). Most self-employed farmers and farm managers were covered in 1954, and additional farmers were covered under liberalized provisions in 1956. Professional self-employment, except dentists, medical professionals, and lawyers, were covered under the 1954 Amendments; except for self-employed doctors, the professional groups excluded in the 1954 Amendments were covered under the 1956 Amendments. Self-employed doctors were covered under the 1965 Amendments. Railroad workers with at least ten years of service were largely covered under the Railroad Retirement Act of 1937, a system parallel and fairly similar to Social Security; those with less than ten years of service were covered by Social Security beginning in 1951. Many employees of the federal government were covered by existing retirement programs in 1935, and therefore excluded from Social Security. The 1950 Amendments extended coverage to a large number of federal employees, mostly part-time or temporary employees, who were not under a Federal retirement system. The 1983 Amendments extended coverage to all civilian employees of the Federal government hired in 1984 onwards. Military employment was covered under the 1956 Amendments (veterans also received noncontributory wage credits for military service between 1940 and 1956).

Our main coding of Social Security coverage is based on the 1950 Census classifications for occupation and industry, which IPUMS uses as a common coding across years (Ruggles et al., 2021), as well as "class of worker" (which includes information on whether a worker was self-employed, a wage or salary earner, or an unpaid family worker). The foregoing discussion noted that there was a regularity of employment condition for coverage in some types of work, such as agricultural wage work and domestic work. A single Census provides information at only one point in time and, in the case of the 1930 Census, does not provide sufficient information on regularity of employment to map into the regularity of employment tests. Hence, our classification implicitly assumes that, for example, a worker who reports his occupation as an agricultural wage worker is employed on a sufficiently regular basis that he would be covered by Social Security in the first

[52] Firefighters and police were treated separately, and do not play a role in our analysis. The 1956 Amendments allowed firefighters and police to be covered in certain states, and the 1967 Amendments allowed firefighters to be covered in all states.

instance (1950 for agricultural wage workers).

We assign 1935 coverage to any wage and salary worker who, based on industry and occupation, was not in one of the original excluded categories. The exclusions were primarily on the basis of industry, although there are a few small categories we exclude on the basis of occupation (such as fishermen and oystermen, or officers and crew members of ships). We assign 1950 coverage to the categories of employment covered on a compulsory basis: agricultural wage workers, domestic workers, and the non-farm self-employed. We also assign 1950 coverage to types of employment that were covered on an elective basis beginning in 1950: nonprofits and state and local government. We treat the following industry codes as implying employment in either nonprofits or state and local government: hospitals, educational institutions, welfare and religious organizations, nonprofit membership organizations, and state and local public administration. Of these, hospitals were one case with a nontrivial number of for-profit entities: based on American Hospital Association (AHA) data, about 30 percent of AHA hospitals were were private in 1935.[53] However, private hospitals tended to be smaller than government or nonprofit hospitals: they held only about 6 percent of hospital beds in 1935, suggesting that they represented only a small share of hospital employment. In the main analysis, we do not attempt to separate out for-profit hospitals, but their presence should tend to attenuate the first stage—which nevertheless is strong—and the size of the reduced form effects. We are aware of no simple way in the Census data of distinguishing nonprofit from government employment in hospitals or education, so we treat them in a parallel way, with elective coverage beginning in 1950. Finally, we categorize police and firefighters separately due to their special coverage under the Social Security Act and exclude them from our analysis. They would not be part of the basis for comparisons anyway, since there would be no variation in Social Security coverage within the same occupation.

If a worker's industry, occupation, or self-employment information is missing or not classified into informative codes, we classify their Social Security coverage status as "unclear" and exclude them from our analysis. Similarly, in some cases a worker's reported industry and occupation may imply different dates of coverage (such as a worker whose occupation is "housekeeper, private household" and industry is "Security and commodity brokerage and investment companies."). These cases we also classify as unclear. In general, we do not exclude, or attempt to correct, apparently inconsistent occupation and industry responses that nevertheless have the same implication for dates of coverage.

To apply our Social Security coverage categories in the OCG sample, which reports occupation and industry according to 1970 Census codes, we first map those codes into 1950 codes, and then apply the Social Security coverage classification described above. Fortunately, despite the changes between 1950 and later years, we still assign a unique Social Security coverage code to each combination of occupation, industry, and self-employment status in the 1970 codes.

A.2.2 Validation of the coverage definition

In addition to the first stage analysis presented in the paper, in Appendix Table A.3 we verify our coverage definition against the 1940 Census, where 5% of the population was asked whether they had obtained a Social Security Number (SSN). The vast majority of those whose 1940 employment is classified as covered in 1935 had obtained an SSN by 1940, while those whose employment is

[53] The AHA data appear in the August issues of *Hospitals: The Journal of the American Hospital Association*. We thank Heidi Williams for generously sharing a digitized version of the data.

classified as not covered until later mostly had not obtained an SSN by 1940. Those whose 1940 employment we do not classify as covered, but who did report having an SSN, may have been employed in a private firm in a largely nonprofit industry; they may have formerly held covered employment; they may have been workers whose employment varied seasonally; or they may have simply registered for an SSN on their own, since coverage of one's employment was not required to get one.

A.3 Pension coverage in the comparison group

Because comparisons are restricted to those within occupation and using variation that arises from exclusions of certain industries from coverage under the original Social Security Act, the comparison group in our empirical analysis is made up largely of certain categories of domestic workers, employees of nonprofits, and employees of state and local governments. To aid interpretation of our main results, it is useful to characterize the prevalence and characteristics of retirement plans for these types of workers. Although we are aware of no systematic information on "retirement plans" for domestic workers (formal or informal), there is enough information on retirement plans for nonprofit and state and local workers to characterize their pension coverage, at least broadly.

A large share of State and local government employees were covered by retirement systems prior to their inclusion in the Social Security system, but within that group, the occupations that form the basis of our main empirical comparisons were relatively less likely to be covered by a retirement system. At the beginning of 1942, when the Census Bureau carried out the first systematic national survey, 46 percent of all State and local government employees were covered by a retirement system; the likelihood of coverage was greater for employees of school systems (58 percent) than for non-school employees (38 percent) (United States Bureau of the Census, 1943, Table 2). This aggregate figure conceals wide disparities across states: in California, New York, Ohio, and Pennsylvania, more than 85 percent of permanent, full time employees were covered, while in a handful of states (Alabama, Idaho, Iowa, Mississippi, Missouri, Nebraska, North Carolina, Oklahoma, Oregon, South Carolina, South Dakota, and Wyoming) less than ten percent of permanent, full-time employees were covered (McCamman, 1943). Coverage under State and local retirement systems grew over the course of the 1940s. In the next nationwide survey, in 1952, the share of State and local government employees covered by a State or a local retirement plan had increased to 67 percent, with a further 9.7 percent covered under Social Security; 74 percent of school employees and 62 percent of non-school employees were covered by State or local systems, with an additional 4.9 percent and 13 percent, respectively, covered by Social Security (United States Bureau of the Census, 1953, Table 2).

Beyond broad categories, there is no systematic information on which types of employees were covered in which localities or states, but the available evidence suggests that the public employees that form part of the basis of our comparison group were, compared to other public employees, less likely to be covered by retirement systems. The State and local employees with the highest rates of coverage by State and local systems were police, firefighters, and teachers (McCamman, 1943; Clark et al., 2003, 2011). Police and firefighters in particular had a long history of coverage and are thought to have had close to universal coverage, typically through localities rather than States, by the end of the 1920s; teachers also had widespread, though far from universal, coverage under either State or local plans (Clark et al., 2003, 2011). We exclude police and firefighters from our analysis entirely, and teachers are used for the comparison only to the extent that they appear

in covered industries (private, for-profit entities) as well as uncovered ones (public or nonprofit organizations).

On the other hand, particularly as of the early 1940s, less than a third of states had general plans for nonschool employees (United States Bureau of the Census, 1943; McCamman, 1953). Altmeyer (1945) noted that within retirement systems for school department employees, it was also often the case as of the mid-1940s that only teachers were eligible, not custodial or nonteaching employees. Consistent with this claim, a 1946 National Education Association study that described characteristics of retirement plans for school employees in 45 states indicated that all 45 states had plans that covered instructional staff of public school systems, but only 15 of these had plans in which custodial employees of school districts could have membership, and 27 had plans in which clerical employees were covered (National Education Association, 1946, p. 18). Retirement plans for public employees grew significantly over the 1940s, which likely reduced this gap to some extent; McCamman (1953) reports that 32 states had state-wide systems for nonschool city employees as of 1950, 21 of which had been introduced since the 1942 survey. Nevertheless, it appears that many of the public-sector workers who form the basis of our comparison group (such as custodial employees of school districts) were not as likely to be covered by retirement systems as were public employees as a whole.

The nature of benefits under State and local retirement plans varied widely, but there are a few generalizations that are relevant to considering how existing retirement coverage in the control group would affect the interpretation of our results. First, public retirement systems tended to have less generous survivor's insurance than Social Security. Most relevant to our analysis, many of these plans had no provision for widow's benefits, or required that retirees accept a lower benefit during their lifetime in order to continue a benefit to a survivor if the retiree died first (McCamman, 1953). This is particularly relevant to our analysis in that co-residence, and perhaps other types of transfers as well, were especially common among widows. Second, benefits under public retirement systems tended to favor lower-paid workers to a lesser extent that Social Security did; they also required longer periods of service than Social Security, and unlike Social Security after the 1939 Amendments, heavily weighted benefits in favor of those with longer periods of service (McCamman, 1953).

Retirement plans for employees of nonprofit organizations existed by 1940, but appear to have been less common than for public employees. Altmeyer (1944) noted that there were voluntary plans made available by professional associations as of the early 1940s, but take-up tended to be low. Based on a survey carried out by the American Hospital Association (AHA), Hayhow (1940) estimated that about 8 percent of nonprofit hospitals in the United States and Canada had pension plans. In 1946, the AHA introduced its own pension plan, which all member nonprofit hospitals were eligible to join (Bugbee, 1946). By early 1949, 120 hospitals and allied organizations had joined (American Hospital Association, 1949); this take-up was relatively modest compared to the 3,160 nonprofit institutions at the time. Like many public retirement systems, coverage under the AHA plan differed from Social Security; of particular note is that securing continuing payments to a retiree's widow after his death required election of a smaller benefit during one's retirement.

Private institutions of higher education comprise another major group of nonprofit organizations. In the 1940s, about one-third of the approximately 1,700 institutions of higher learning were public, and two-thirds were private; of those two-thirds, more than 90 percent were nonprofit organizations (Murray and Smith, 1940; McCamman, 1948). McCamman (1948) reported that about 30 percent of private institutions (or 338 total) had retirement plans, 78 percent of which were

through contracts with the Teachers Insurance and Annuity Association (TIAA). Faculty in these institutions are not part of our empirical comparisons, since they do not appear in both covered and uncovered industries (and because of the rarity of for-profit higher educational institutions, we classify all institutions as public or nonprofit); instead, the employees most relevant to our comparisons are maintenance and other nonprofessional employees. For the most part, contracts that private institutions held with TIAA did not include these types of employees (McCamman, 1948).[54] Of the smaller share that were not covered through TIAA contracts, some were likely affiliated with religious organizations; it is worth noting that Church pension funds frequently covered all employees (Murray and Smith, 1940).

A.4 Estimating the gains from migration based on evidence on *non*-migration gains

In Section 6, we use a variety of approaches to estimate migration's role in driving the financial gains from Social Security to recipients' children. Here, we take a complementary approach based on evidence on the likely *non*-migration gains from Social Security to recipients' children. Such evidence, together with estimates of the overall gain, allow one to infer the gain from migration as a residual.

The main non-migration way in which Social Security might benefit recipients' children is by increasing net downstream transfers from recipients to their children. As is well known, increases in the income or wealth of older individuals tend to trigger offsetting changes in family transfers (e.g., Barro, 1974; Becker, 1974). A relatively high estimate is that a dollar increase in the present value of an older household's income or wealth increases transfers to its children by 25 cents.[55] In addition, Social Security, in displacing family old-age support, might free up some of the time of the children who would otherwise provide such support. As a benchmark, suppose that the corresponding benefit to recipients' children is another 25 cents per dollar of Social Security benefits, for an overall benefit of 50 cents per dollar of Social Security benefits. How much of the overall gain to recipients' children would such gains account for? We estimate that a ten-year difference in the predicted earliest age of father's eligibility—roughly the difference in actual earliest age of eligibility between 1935- and 1950-covered fathers from the 1877 birth cohort—is associated with a roughly $15,000 difference in the expected present value of Social Security benefits. If the total non-migration benefit to recipients' children was 50 cents per dollar, the resulting increase in transfers would have a present value of about $7,500, which is about one-fifth of the estimated overall gain. So this approach suggests that migration gains account for about four-fifths of the total gains.

A.5 Potential efficiency gains from Social Security smoothing the costs of old-age support

Social Security may have facilitated the exploitation of gains from trade in old-age support that were otherwise foregone due to missing markets and transaction costs. Social Security likely smoothed the financial component of old-age support across individuals and over time, in a way that would have been hard to achieve otherwise. Such a smoothing, in turn, likely reduced the

[54] Similar to the other types of retirement systems described here, at the time of retirement, TIAA annuity contracts offered the choice between an annuity that ceased at the retiree's death and one that made smaller payments that continued until his widow's death, if the retiree died first.

[55] For example, Mukherjee (forthcoming) estimates a pass-through rate closer to 15%.

opportunity cost of migrating away from one's parents. In this way, Social Security may have facilitated a more efficient allocation of people across places. Evidence in keeping with this idea includes that Social Security appears to have significantly increased the average total family earnings of recipients and their children.

Social Security likely smoothed the financial component of old-age support across individuals and over time. Without Social Security, the costs of old-age support were highly concentrated both across individuals and over time, often borne primarily by one or two working-age children during a relatively brief period of time. This concentration likely arose from a combination of missing markets and key features of old-age and family old-age support technologies. Missing markets, especially the lack of insurance markets for hedging the risk in how much old-age support will be demanded of a particular individual, likely greatly limited the extent to which individuals could insure old-age support costs and so the extent to which such costs were spread across individuals, especially across families but also within families. Capital market imperfections likely limited the extent to which individuals could smooth old-age support costs over time, especially by limiting the extent to which they could finance costs out of future income by borrowing. Key features of old-age and family old-age support technologies likely further concentrated old-age support costs. Old age is a relatively brief portion of the life cycle. There are large economies of scale in providing old-age support in the form of shared housing or shared location, which not only must occur at the same time that the older generation receives support but also are likely most efficiently provided by one or few providers (so that others need not live in the same house or location, and to economize on coordination costs). As a result, without Social Security the equilibrium distribution of old-age support costs was highly concentrated, most often involving one child living with the parents and almost always involving at least one child living near the parents, even at the expense of being poorly matched to the local labor market. The first-best allocation, by contrast, would likely involve much less concentration both across individuals and over time, since that would involve both more risk-sharing and a better matching of individuals to labor markets. Hence, the equilibrium without Social Security may have been highly inefficient; there may have been large foregone gains from trade in the provision of old-age support due to missing markets and transactions costs.[56] Social Security likely made the financial component of old-age support much more uniform, since it involved virtually all working-age individuals paying a modest share of their earnings as Social Security taxes throughout their working lives, independent of the level and timing of their parents' demand for and their siblings' supply of old-age support.

Social Security, in smoothing the financial component of old-age support across individuals and over time, likely reduced the opportunity cost of migration for recipients' children. Without Social Security, the highly concentrated nature of family old-age support costs likely created a strong incentive for providers to live with or near their parents to economize on the large transfers they were making during their parents old age. Much of family old-age support takes the form of in-kind transfers of shared housing, help with chores, or personal care, which require the provider to be in the same location as the parent. In this sense, intensive family old-age support is similar to other jobs in that "workers" live in the same labor market in which they work. Such scale

[56]Our evidence of the large, migration-driven gains to recipients' children from Social Security is suggestive of there having been a large gap between the equilibrium without Social Security and the first best. Even upper-bound estimates of the increase in the cost of living from families switching from shared to separate locations in response to greater Social Security coverage are much smaller than the associated migration gains. That co-residence appears to be an inferior good reinforces this conclusion that the equilibrium without Social Security was highly inefficient.

economies likely comprised a large opportunity cost of migrating away from one's parents, since doing so would either preclude living with or near one's parents or require an additional move to do so. Social Security, in smoothing the financial component of old-age support, likely reduces the costs of the would-be primary providers of financial support. This reduces their gains from living near their parents during their parents' old age and so reduces their opportunity cost of migration. In this way, Social Security likely relaxes for at least some would-be support-providing children a "location constraint" that they live near their elderly parents. This could potentially lead them to sort to labor markets that better match their skills.[57]

This discussion suggests that Social Security may have increased efficiency, helping people realize gains from trade in old-age support that were previously foregone due to missing markets and transactions costs. Social Security likely smoothed the financial component of old-age support both across individuals and over time, transforming them from being highly concentrated (often borne mainly by just one or two children for only a brief spell during their parents' old age) to more disperse (borne partly in the form of modest tax payments by all workers throughout all of their working years). This transformation likely brought the equilibrium allocation closer to the first-best allocation and so increased efficiency.[58] In other words, Social Security may have helped people realize gains from insuring the financial component of old-age support and smoothing it over time that were previously limited by incomplete markets for insuring old-age support costs and smoothing such costs over time. The absence or imperfection of such markets likely meant that without Social Security, the costs of old-age support were much larger than in the first-best allocation with complete markets, largely but not only because many individuals forewent much better earning opportunities in other locations in order to live near or with their parents to support them in their old age.[59]

To fix ideas, consider the thought experiment of offering someone behind the veil of ignorance the choice between (i) the no-Social Security allocation of old-age support, which is highly-concentrated with a small share of adult children providing most of the support, and (ii) a Social Security-like distribution of old-age support, which smooths the financial component of old-age support (though likely not the personal-care component) both across individuals and over time. The idea is that from this original position behind the veil of ignorance, one would likely prefer (ii). This is partly because it partially insures risk in how much old-age support will be demanded of one and partly because it significantly reduces the average cost of providing such support, by reducing the demand for living with or near one's parents in order to provide such support intensively. Of course, there is no strict requirement that the no-Social Security allocation of old-age support be so highly-concentrated and involve so many adult children living with or near their elderly parents. Instead, it arises as an equilibrium given the constraints imposed by incomplete markets and transaction costs in smoothing old-age support costs across individuals and over time

[57] Another important missing market might have been that for insuring the risky returns to migration. Migration is costly and risky, so even if it pays off handsomely on average, it may not pay off enough to finance old-age support in certain states of the world.

[58] The idea is that smoothing across individuals (insurance from behind the veil) and smoothing over time (life cycle smoothing) tend to be valuable given diminishing marginal utility, and that without Social Security such smoothing was greatly limited by incomplete markets for insuring old-age support costs and imperfect capital markets.

[59] In addition to incomplete markets for insurance of old-age support costs (i.e., ex ante trade in old-age support costs across individuals) and imperfect capital markets for smoothing over time, other factors that may have limited families' ability to minimize old-age support costs include incomplete markets for insuring the gains from migration and perhaps underestimation of the gains from migration.

without Social Security and perhaps due to other factors as well. In this thought experiment, Social Security enables families to smooth the financial component of old-age support, which reduces the opportunity cost of migrating for some of recipients' children, which may facilitate their migration to better-matched labor markets.[60]

Evidence in keeping with the idea that Social Security facilitated a more efficient allocation of people across places includes that Social Security appears to have significantly increased the average total family earnings of recipients and their children. This raises the question: Was the main barrier discussed above—incomplete markets—enough to explain why families were forgoing large gains in family income on average without Social Security (but not with Social Security)? Answers to this question are necessarily speculative given a lack of exogenous variation in key factors that would enable direct empirical tests. Nevertheless, we consider some other factors that may have contributed to leading families to forgo large gains in family income on average without Social Security.

- Misperception of gains from migration: One possibility is that families were forgoing large gains without Social Security not because such gains were worth forgoing in order to facilitate family old-age support but because they did not realize that the gains from migration were so large. In this view, Social Security, in reducing the demand for family old-age support, led some families to migrate and enjoy large gains—much bigger than they anticipated. This seems not unlikely as a potential contributor. Given that gains from migration have been found to be quite large in many contexts, it seems plausible that families might not have realized that they were so large. To the extent that families were migrating less than would have been efficient without Social Security, Social Security-induced increases in migration have direct efficiency consequences even beyond the fiscal externalities from higher earnings and so net tax payments (since it has a first-order effect on the families themselves).[61]

- Norms or transaction costs within families: A variety of norms or transaction costs within families may have limited families' ability to exploit potential gains from trade. For example, norms might have favored local, in-kind old-age support over cash old-age support, such as a norm that without Social Security at least one child stays near their parents (whereas with Social Security children can move where they would like). A "transaction cost" reason for such a norm of shared location might be that non-local children will not provide much old-age support, perhaps due to lower connection with their parents. While there is a clear force toward norms that promote efficiency, to mean anything norms must be at least somewhat durable and a norm that promoted efficiency at one time may not continue to do so

[60] Note that Social Security did not *force* families to smooth the costs of old-age support; the average family retained their no-Social Security allocation in their with-Social Security constraint. If families fully exploit within-family gains from trade (i.e., efficient families), then the fact that families appear to have chosen a smoother distribution of the financial component of old-age support with Social Security than without is suggestive that they prefer such a distribution, i.e., that such a distribution is more efficient. In other words, it suggests that with complete markets people would have chosen a distribution of old-age support closer to that under Social Security than that without Social Security, i.e., that Social Security increased efficiency, i.e., that Social Security effectively created valuable otherwise-missing markets for trade in old-age support.

[61] Underestimation of gains from migration could potentially help explain other findings in the literature as well. For example, a variety of quasi-experimental and structural studies suggest that migration confers significant financial gains on those induced into migrating by various changes in the environment. A leading explanation is that migration costs are very large. A complementary explanation is that individuals might underestimate migration benefits.

indefinitely. It seems not unlikely that such norms or transaction costs contributed to why families forewent large income gains on average without Social Security.

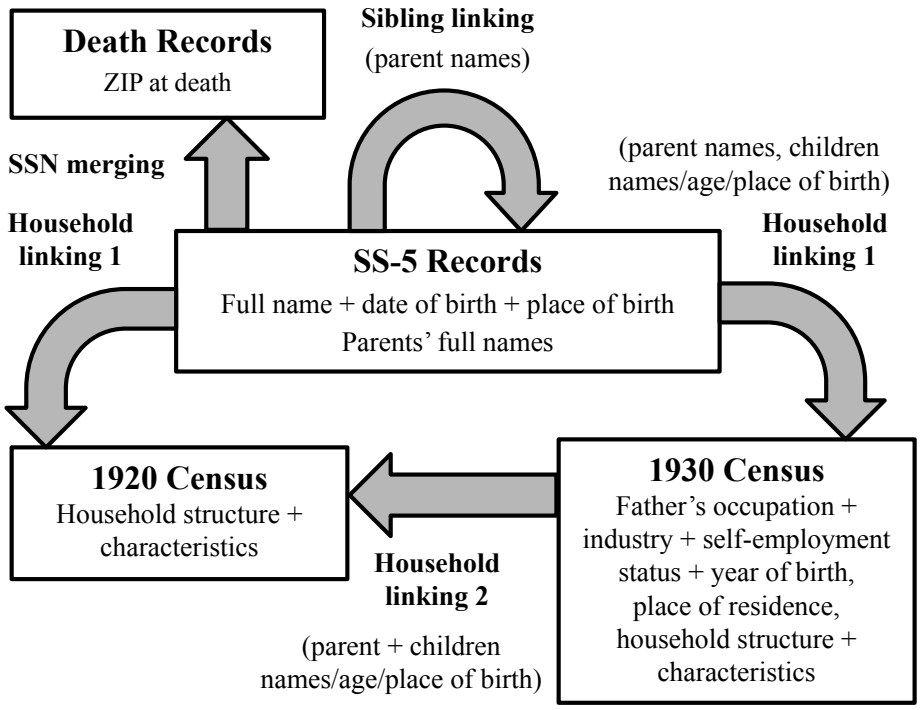
- Shift from non-market to market time: Social Security, in displacing in-kind family old-age support, may have "freed up" some time of would-be support-providers, which may have led them to increase their labor supply and earnings. While this seems likely to have been a contributing factor, both that the financial gains to recipients' children appears to have exceeded the associated Social Security benefits and that the effects are concentrated on sons rather than daughters cast doubt on this mechanism. It seems unlikely that this mechanism would create gains that were so large on average and concentrated on sons, since sons' labor hours were likely less affected.

- Income effects: Social Security conferred large net windfalls on the families of early recipients, and independent living—and possibly living in different locations—is a normal good, so Social Security reduced the demand for shared location. While income effects could in principle explain a shift toward independent living, such a mechanism seems hard to square with the finding that Social Security likely increased total family earnings. If without Social Security families were too poor to live independently or move, would not they also be too poor to forgo the large average income gains from migrating? In fact, given the large gains from migration, one might expect income effects to if anything *reduce* migration. The idea is that indulging whatever preferences lead to forgoing a lot of income on average is presumably a normal good, so Social Security, in conferring windfalls on early recipients, would tend to increase the extent to which families choose to indulge such preferences, at the expense of earning lower income.

Of course, we emphasize that such considerations are highly speculative given the lack of direct evidence available. The extent to which any efficiency gain from Social Security-induced migration came from substituting for missing markets versus correcting mistakes does not affect the size of the gain but is potentially relevant for understanding its generalizability to other contexts (e.g., contexts in which either the key markets or beliefs differ from our context). This is an interesting question for future research.

Appendix Figures and Tables

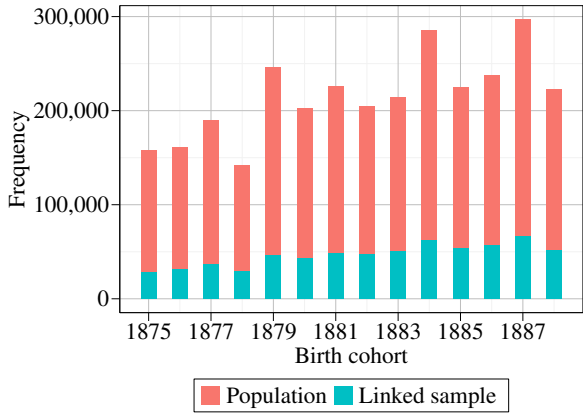
Figure A.1: Overview of Numident death-census linked sample construction



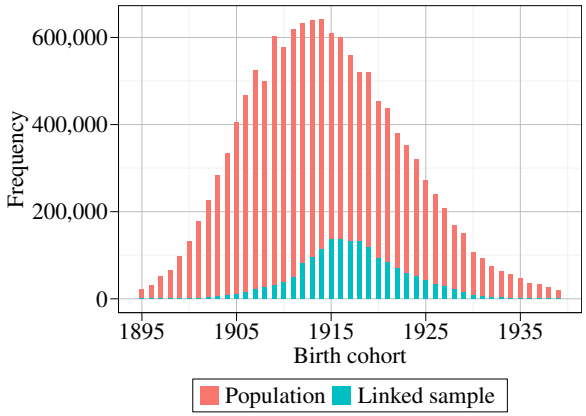
Notes: The key variables from each dataset are indicated in the boxes. The variables used for each type of linkage are indicated in parenthesis.

Figure A.2: Coverage of men born in 1875–1888 and their children born in 1895–1939

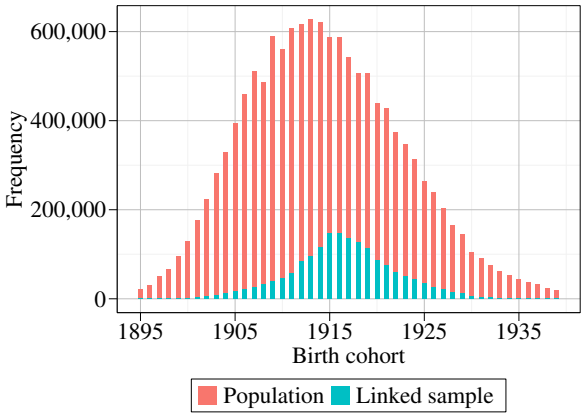
(a) Men born in 1875–1888, covered by Social Security in 1935 or 1950 based on their 1930 employment information



(b) Sons of men born in 1875–1888

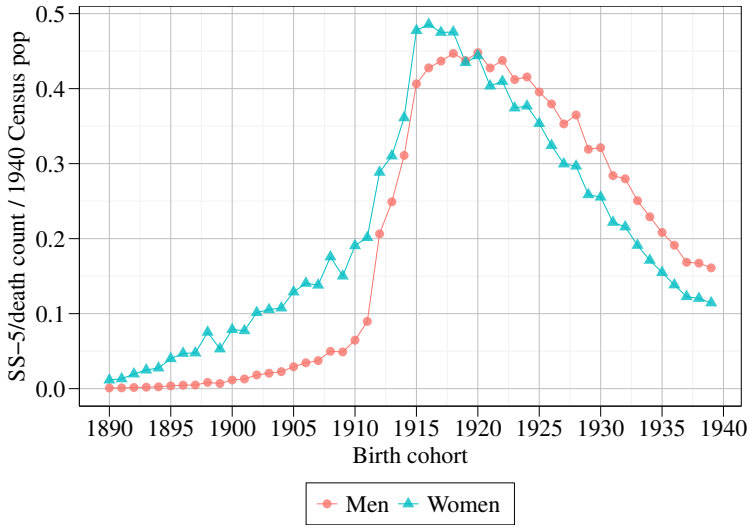


(c) Daughters of men born in 1875–1888



Notes: Population in panel (a): men born in 1875–1888 in the 1930 Census, whose employment was covered by Social Security in 1935 or 1950 based on their employment information in 1930. Linked sample in panel (a): fathers in linked sample with analogous restrictions. Population in panels (b) and (c): sons/daughters born in 1895–1909/1910–1919/1920–1929/1930–1939 in the 1910/1920/1930/1940 Censuses, whose fathers were born in 1875–1888. Linked sample in panels (b) and (c): sons/daughters born in 1895–1939 in linked sample, whose fathers were born in 1875–1888.

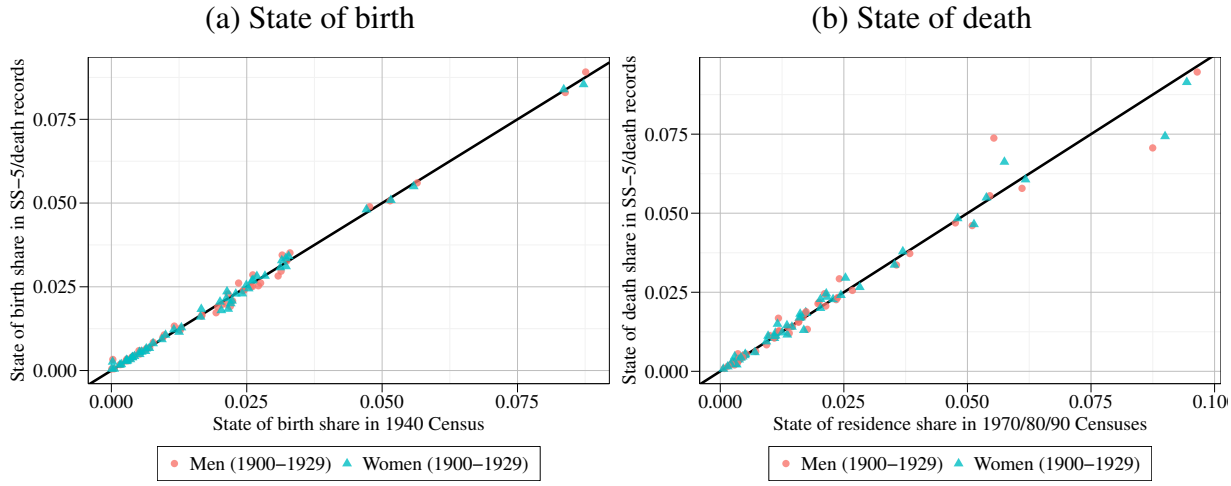
Figure A.3: Cohort coverage of SS-5/death records by sex (1890–1939 cohorts)



Notes: This figure plots the implied cohort coverage rate of SSNs that can be found in both SS-5 and death records with non-missing death ZIP information, defined as the number of SSNs divided by the corresponding number of individuals in the 1940 Census, separately by gender (legend). It therefore captures the share of individuals who survived until 1940 who appear in Numident records.

Source: Numident SS-5 and death records, 1940 US Census.

Figure A.4: State of birth/death coverage of SS-5/death records by sex (1900–1929 cohorts)

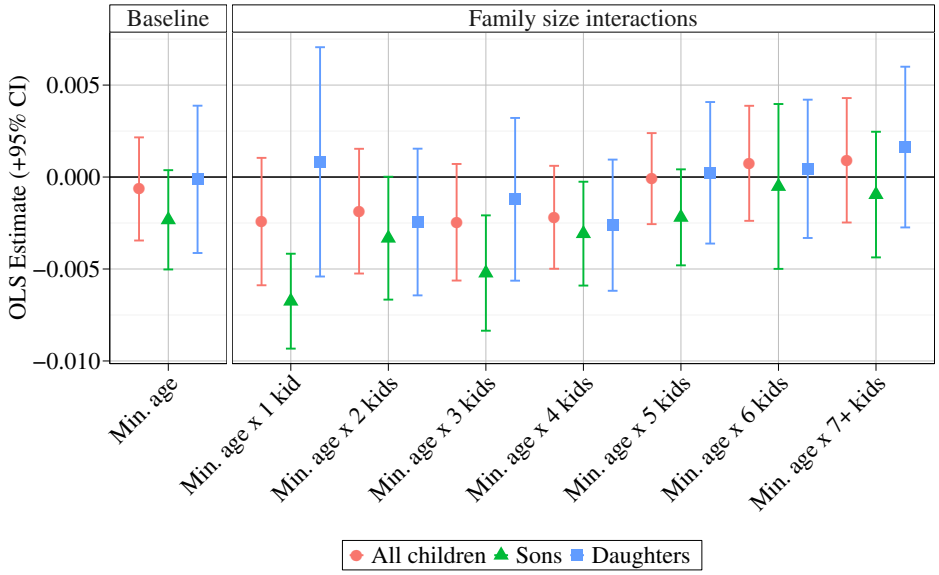


Notes: Panel (a) plots state of birth shares among SSNs that can be found in both SS-5 and death records with non-missing death ZIP information against the corresponding shares in the 1940 Census for individuals born between 1900 and 1929, separately by gender. For the same cohorts and set of SSNs, Panel (b) plots state of death shares in death records against state of residence shares in the 1970 Census (1900–1909 cohorts), 1980 Census (1910–1919 cohorts), and 1990 Census (1920–1929 cohorts), separately by gender.

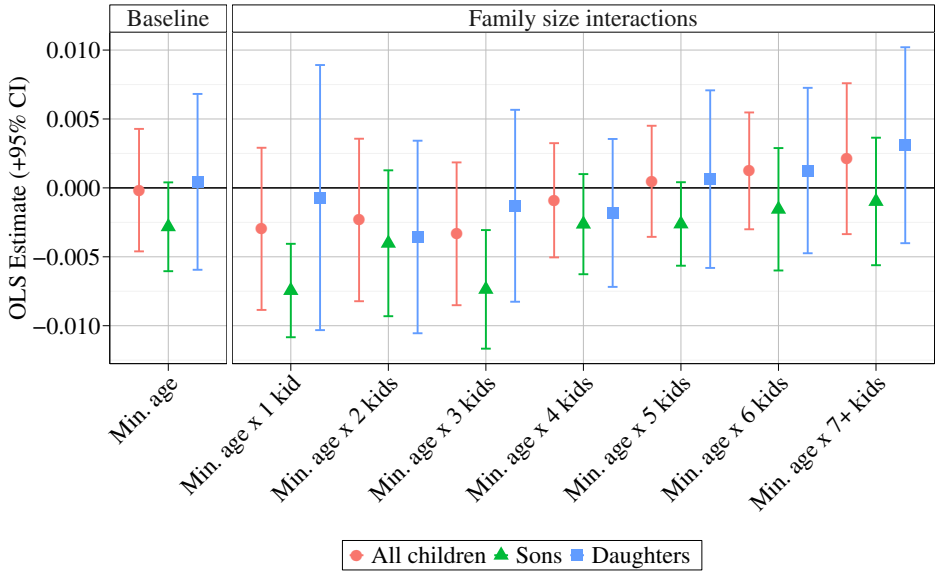
Source: Numident SS-5 and death records, 1940, 1970, 1980 and 1990 US Censuses.

Figure A.5: The impact of father's later Social Security coverage on children's death ZIP outcomes (in logs)

(a) Death ZIP log mean AGI in 2001



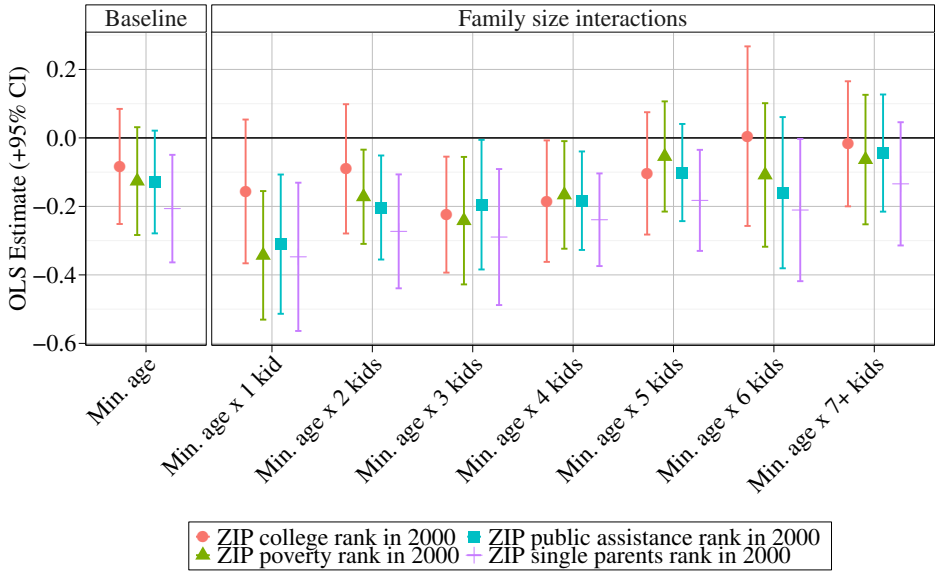
(b) Death ZIP log median house values in 2000



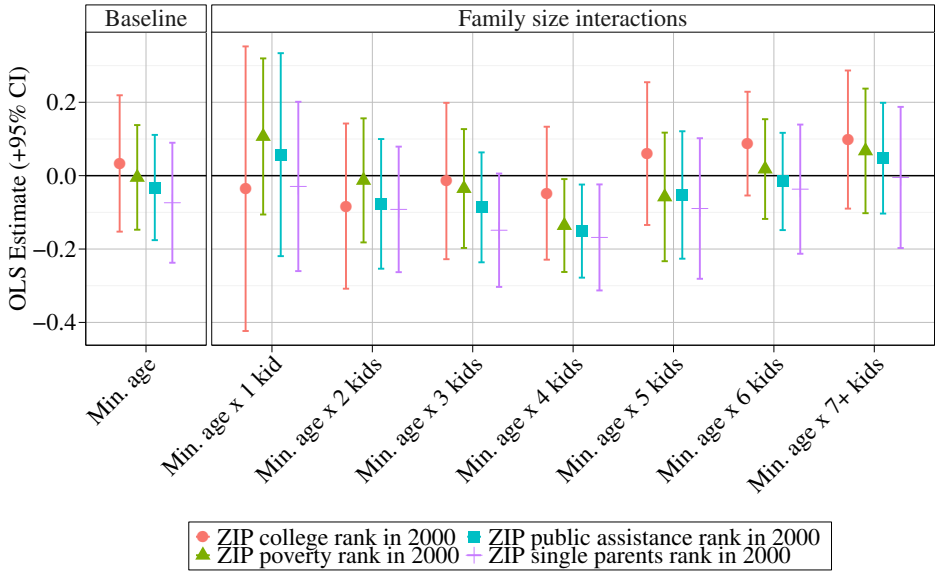
Notes: Figures display coefficients on father's minimum age of Social Security eligibility from equation (2), where the outcome is the log of mean AGI or median house values in his children's ZIP code at death (panel title). Baseline estimates shown in the left panel of each graph; estimates interacting minimum age with fixed effects for number of children shown in the right panel. Colors/shapes indicate separate estimates for all children, sons only, or daughters only (legend). All regressions include fixed effects for father's occupation × year of birth, Social Security coverage category, race, and state/country of birth. Observations weighted using IPWs. 95% confidence intervals based on standard errors clustered at the level of father's industry of employment.

Figure A.6: The impact of father's later Social Security coverage on children's death ZIP rank in terms of 2000 economic indicators

(a) Sons



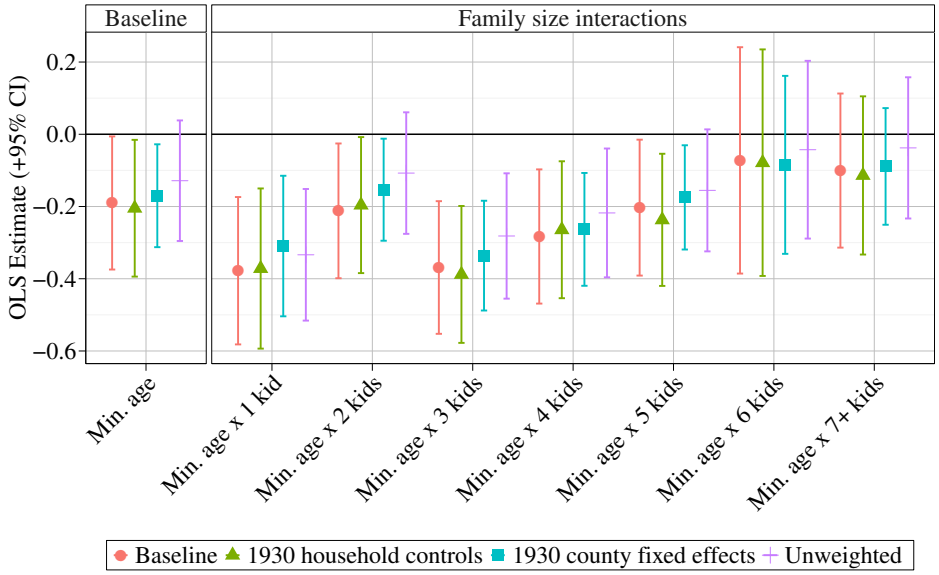
(b) Daughters



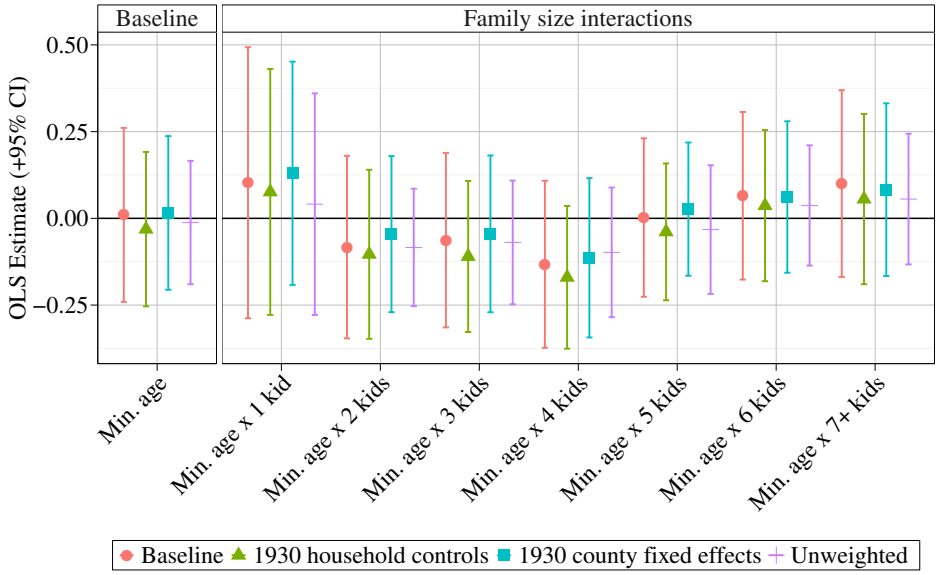
Notes: Figures display coefficients on father's minimum age of Social Security eligibility from equation (2), where the outcome is the average percentile rank of his children's ZIP code at death in terms of various SES measures (panel title). Baseline estimates shown in the left panel of each graph; estimates interacting minimum age with fixed effects for number of children shown in the right panel. Colors/shapes indicate separate estimates for all children, sons only, or daughters only (legend). All regressions include fixed effects for father's occupation × year of birth, Social Security coverage category, race, and state/country of birth. Observations weighted using IPWs. 95% confidence intervals based on standard errors clustered at the level of father's industry of employment.

Figure A.7: The impact of father's later Social Security coverage on children's death ZIP rank in terms of 2001 AGI: Robustness checks

(a) Sons

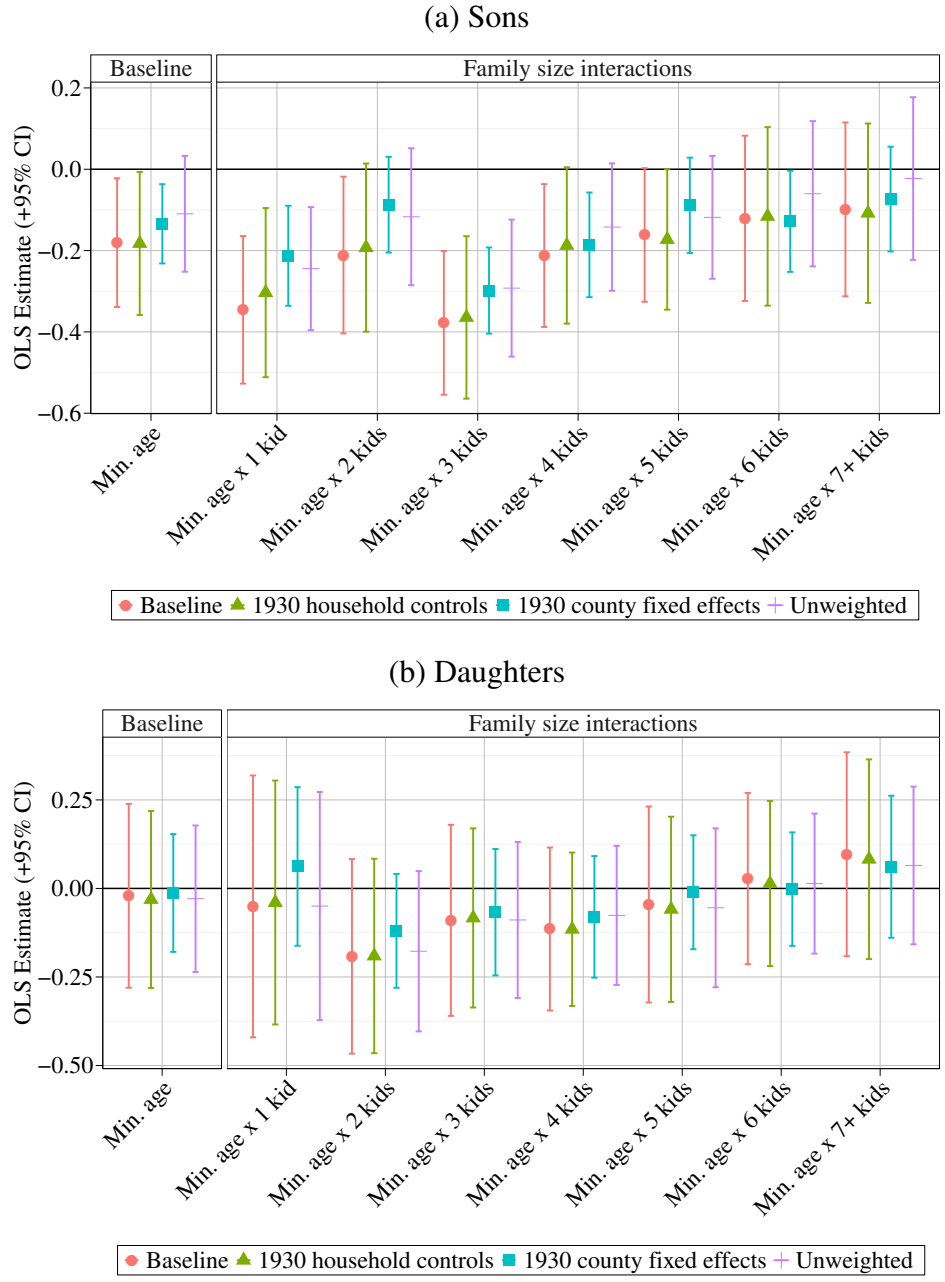


(b) Daughters



Notes: Figures display coefficients on father's minimum age of Social Security eligibility from equation (2), where the outcome is the average percentile rank of his children's ZIP code at death in terms of mean AGI, either for sons or daughters (panel title). Baseline estimates shown in the left panel of each graph; estimates interacting minimum age with fixed effects for number of children shown in the right panel. Colors/shapes indicate alternative specifications (legend). All regressions include fixed effects for father's occupation × year of birth, Social Security coverage category, race, and state/country of birth. Observations weighted using IPWs (except unweighted specification). 95% confidence intervals based on standard errors clustered at the level of father's industry of employment.

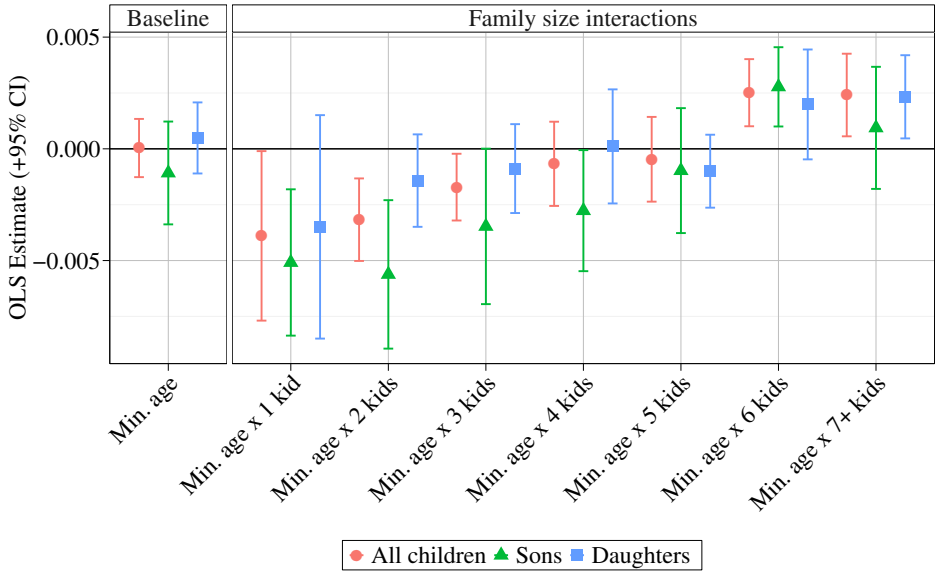
Figure A.8: The impact of father's later Social Security coverage on children's death ZIP rank in terms of 2000 house values: Robustness checks



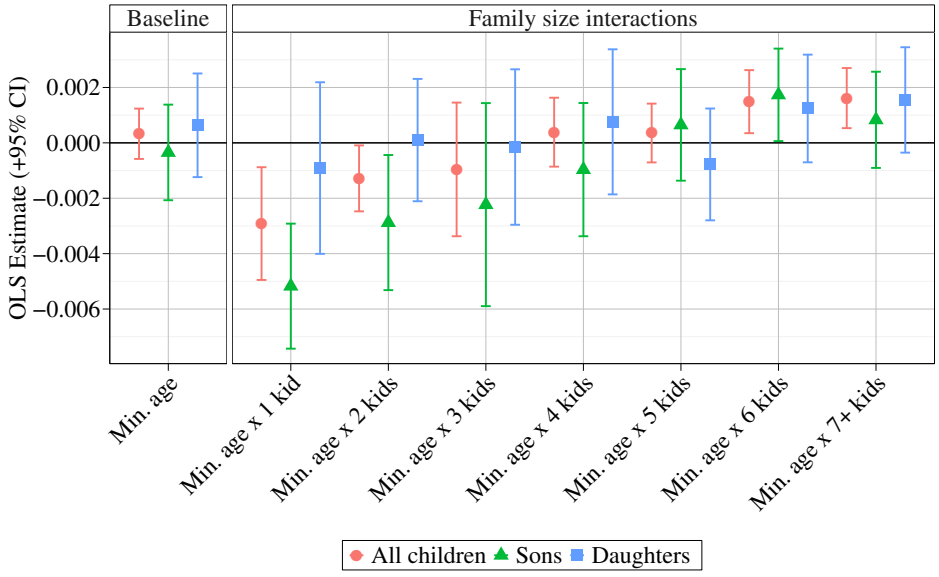
Notes: Figures display coefficients on father's minimum age of Social Security eligibility from equation (2), where the outcome is the average percentile rank of his children's ZIP code at death in terms of median house values, either for sons or daughters (panel title). Baseline estimates shown in the left panel of each graph; estimates interacting minimum age with fixed effects for number of children shown in the right panel. Colors/shapes indicate alternative specifications (legend). All regressions include fixed effects for father's occupation × year of birth, Social Security coverage category, race, and state/country of birth. Observations weighted using IPWs (except unweighted specification). 95% confidence intervals based on standard errors clustered at the level of father's industry of employment.

Figure A.9: The impact of father's later Social Security coverage on children's long-run migration: Alternative migration outcomes

(a) Died outside parents' 1930 state of residence

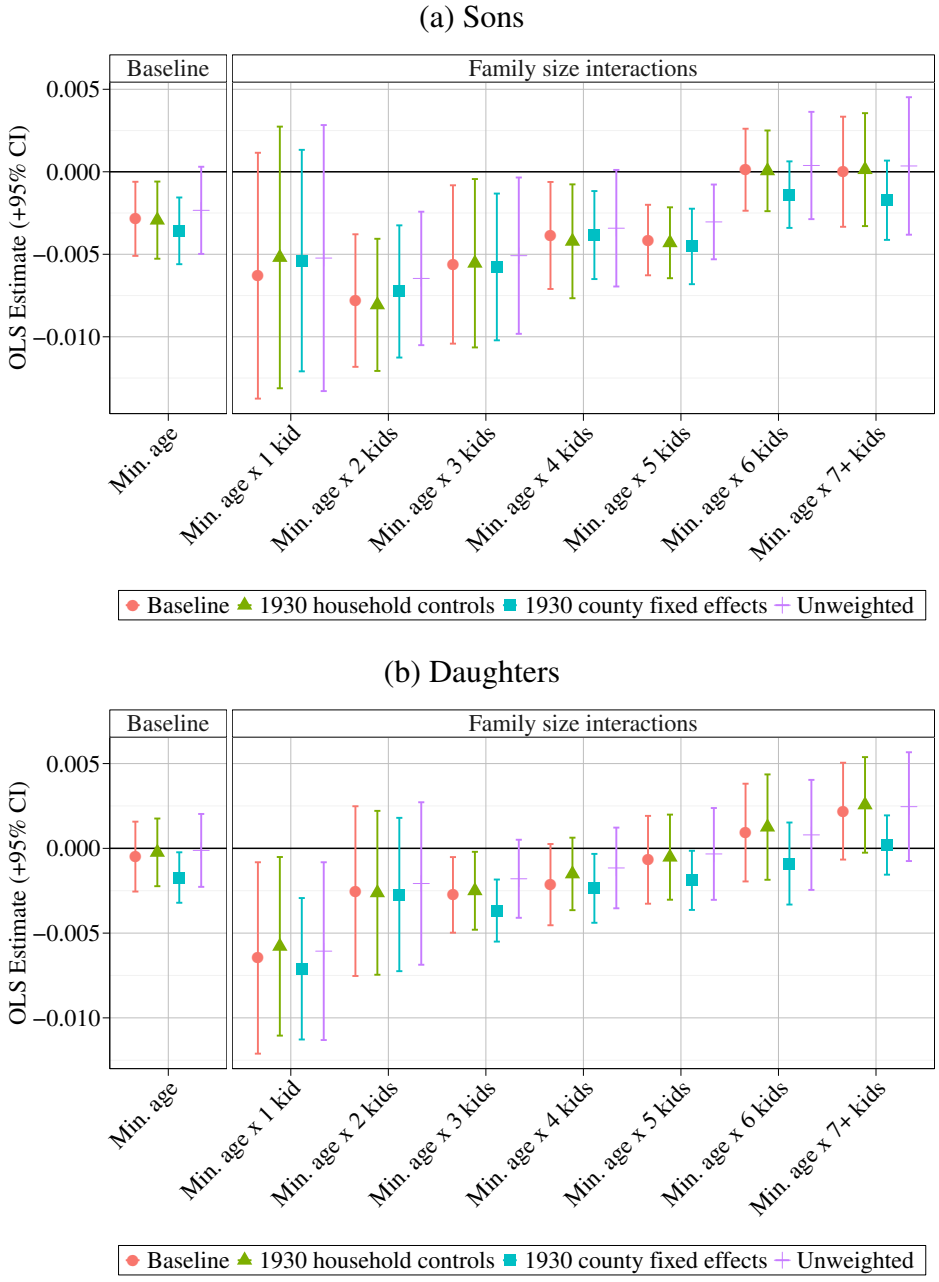


(b) Died outside parents' 1930 Census region of residence



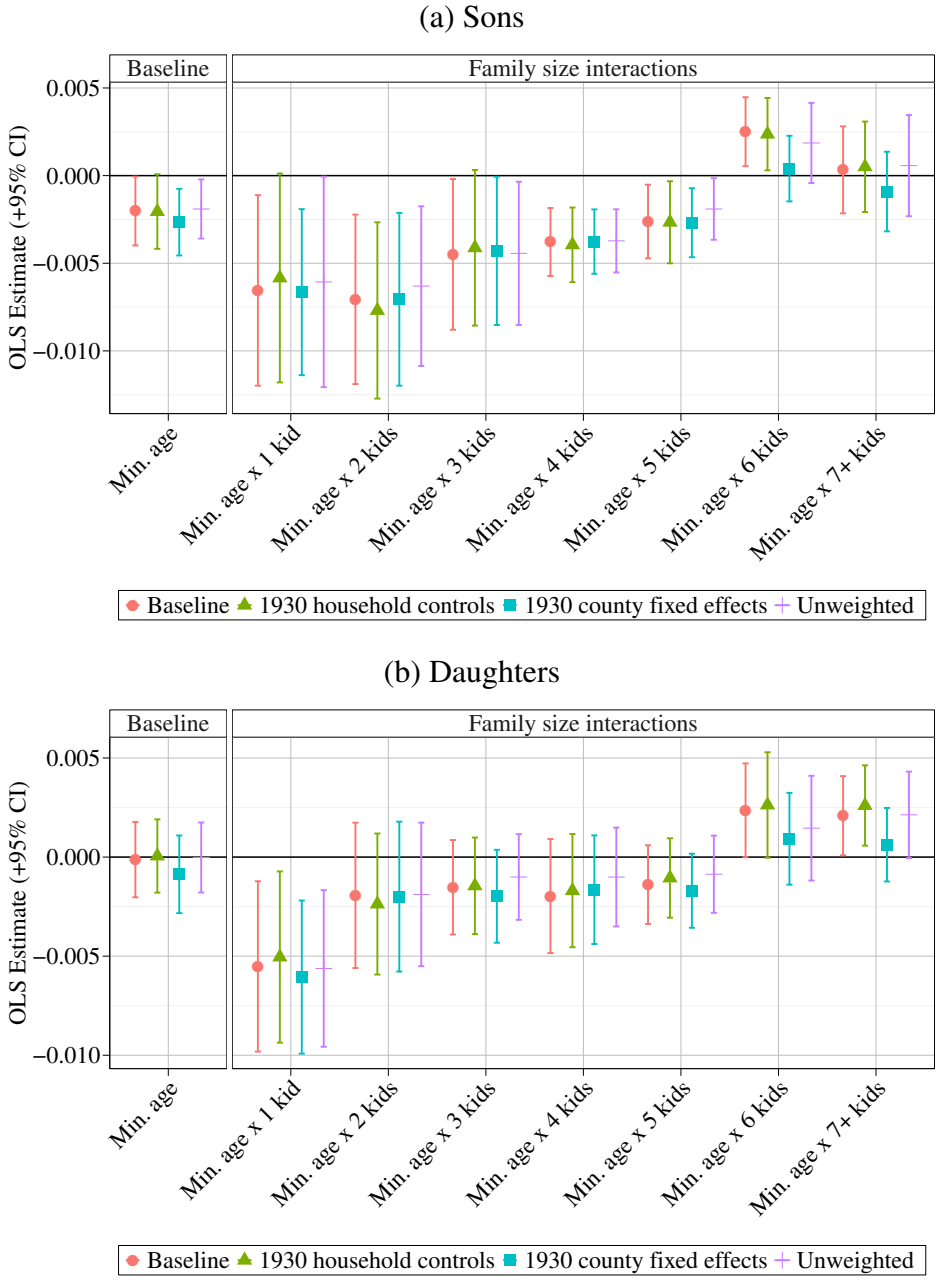
Notes: Figures display coefficients on father's minimum age of Social Security eligibility from equation (2), where the outcome is the share of his children who died outside his 1930 state or Census region of residence (panel title). Baseline estimates shown in the left panel of each graph; estimates interacting minimum age with fixed effects for number of children shown in the right panel. Colors/shapes indicate separate estimates for all children, sons only, or daughters only (legend). All regressions include fixed effects for father's occupation × year of birth, Social Security coverage category, race, and state/country of birth. Observations weighted using IPWs. 95% confidence intervals based on standard errors clustered at the level of father's industry of employment.

Figure A.10: The impact of father's later Social Security coverage on children's probability of dying outside their parents' 1930 county of residence: Robustness checks



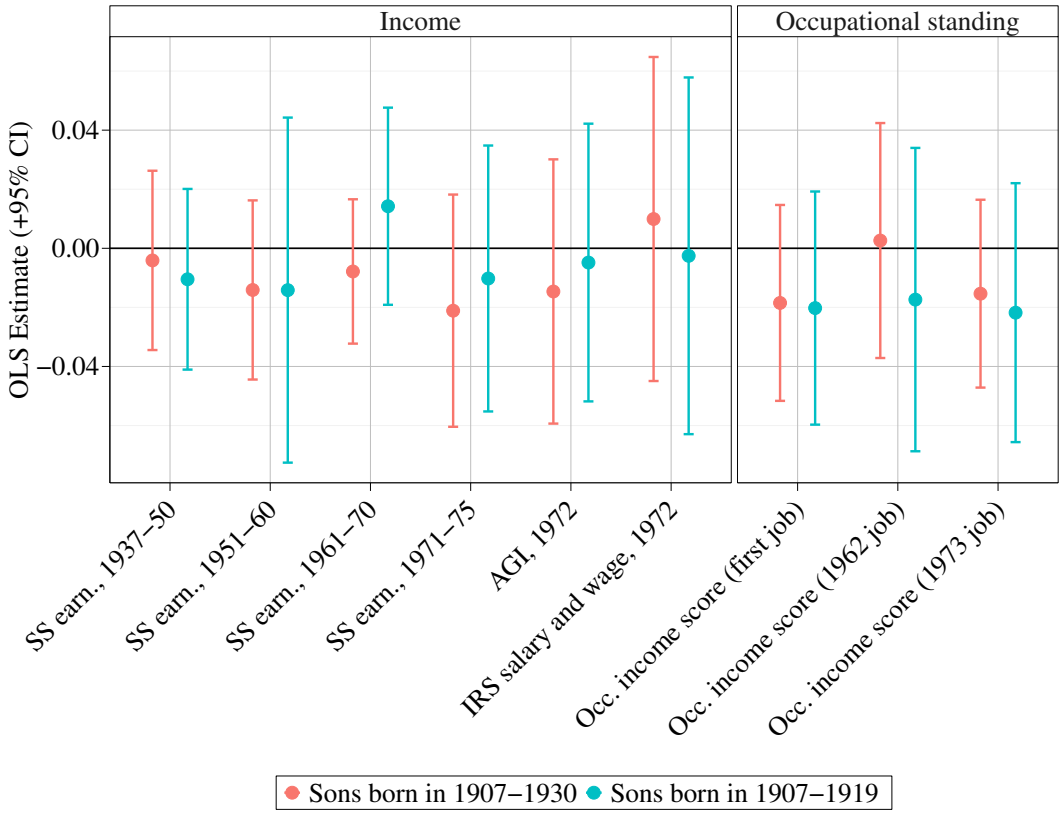
Notes: Figures display coefficients on father's minimum age of Social Security eligibility from equation (2), where the outcome is the share of his children who died outside his 1930 county of residence, either for sons or daughters (panel title). Baseline estimates shown in the left panel of each graph; estimates interacting minimum age with fixed effects for number of children shown in the right panel. Colors/shapes indicate alternative specifications (legend). All regressions include fixed effects for father's occupation × year of birth, Social Security coverage category, race, and state/country of birth. Observations weighted using IPWs (except unweighted specification). 95% confidence intervals based on standard errors clustered at the level of father's industry of employment.

Figure A.11: The impact of father's later Social Security coverage on children's probability of dying outside their parents' 1930 CZ of residence: Robustness checks



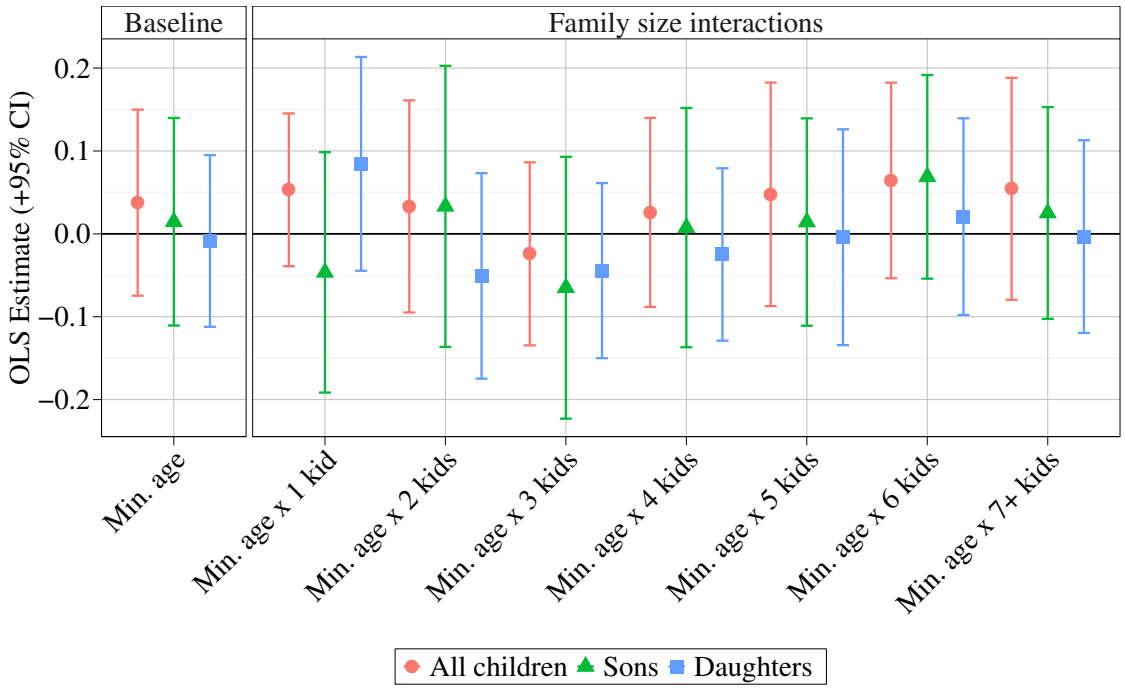
Notes: Figures display coefficients on father's minimum age of Social Security eligibility from equation (2), where the outcome is the share of his children who died outside his 1930 commuting zone of residence, either for sons or daughters (panel title). Baseline estimates shown in the left panel of each graph; estimates interacting minimum age with fixed effects for number of children shown in the right panel. Colors/shapes indicate alternative specifications (legend). All regressions include fixed effects for father's occupation × year of birth, Social Security coverage category, race, and state/country of birth. Observations weighted using IPWs (except unweighted specification). 95% confidence intervals based on standard errors clustered at the level of father's industry of employment.

Figure A.12: The impact of father's later Social Security coverage on sons' probability of having non-missing positive income and a non-missing occupation



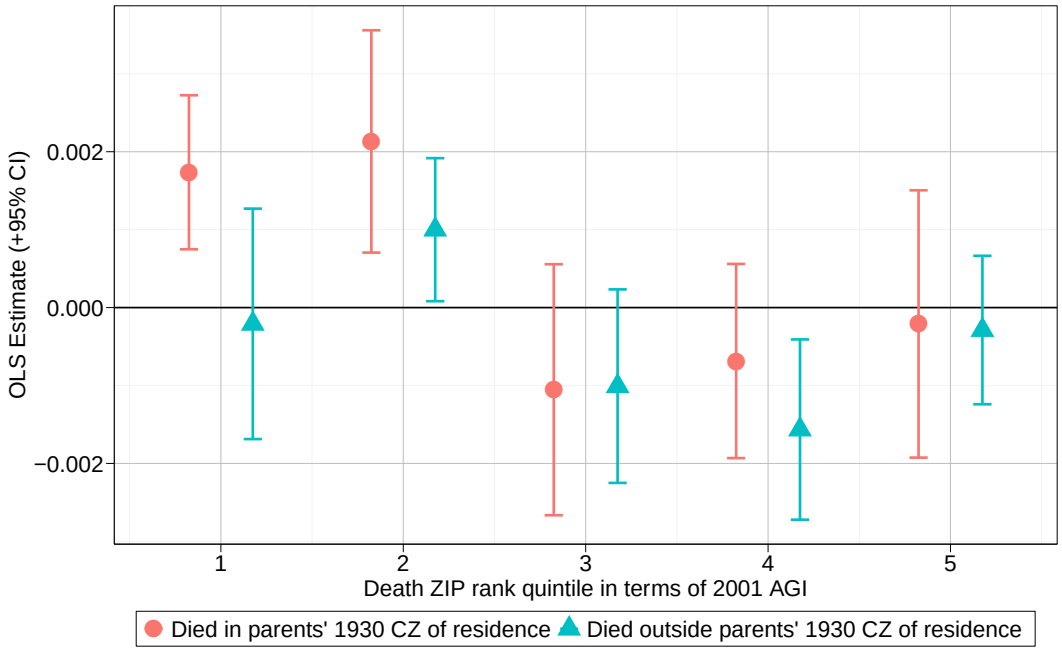
Notes: Figure displays coefficients on father's minimum age of Social Security eligibility from equation (2), where the dependent variable is an indicator for a specific outcome (x-axis) being non-missing and positive (left panel), or simply non-missing (right panel). Colors/shapes indicate separate estimates for sons born in 1907–1930 and sons born in 1907–1919 (legend). All regressions include fixed effects for father's occupation $\times$ year of birth, Social Security coverage category, and son's year of birth, race, and state/country of birth. Observations weighted using OCG sampling weights. 95% confidence intervals based on standard errors clustered at the level of father's industry of employment.

Figure A.13: The impact of father's later Social Security coverage on children's death CZ rank in terms of 2001 AGI



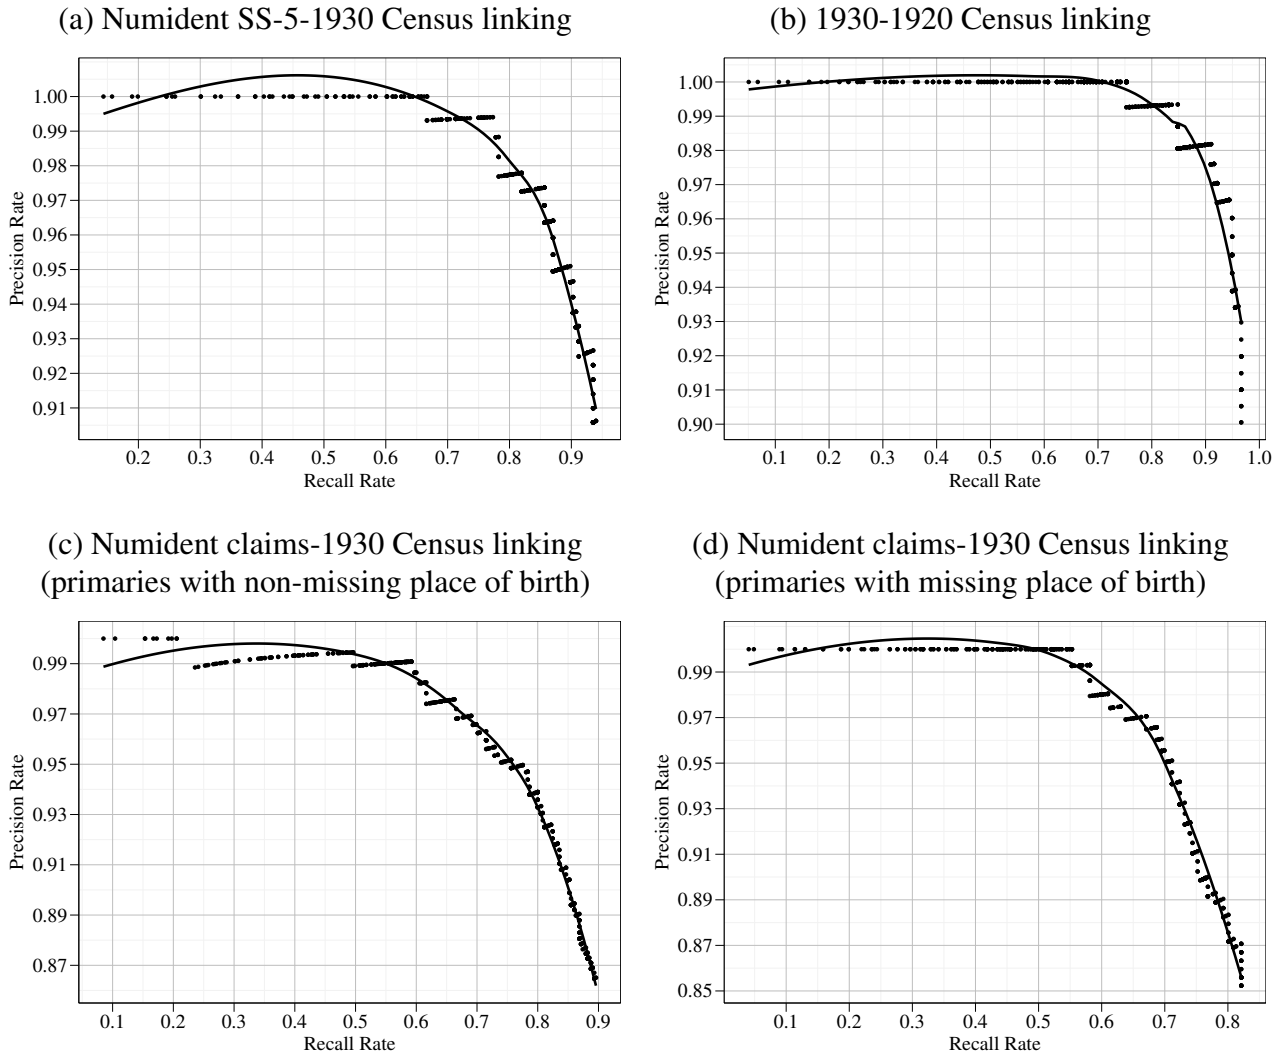
Notes: Figure displays coefficients on father's minimum age of Social Security eligibility from equation (2), where the outcome is the average percentile rank of his children's CZ at death in terms of mean AGI. Baseline estimates shown in the left panel; estimates interacting minimum age with fixed effects for number of children shown in the right panel. Colors/shapes indicate separate estimates for all children, sons only, or daughters only (legend). All regressions include fixed effects for father's occupation × year of birth, Social Security coverage category, race, and state/country of birth. Observations weighted using IPWs. 95% confidence intervals based on standard errors clustered at the level of father's industry of employment.

Figure A.14: The impact of father's later Social Security coverage on sons' joint death ZIP rank quintile and migration status



Notes: Figure displays coefficients on father's minimum age of Social Security eligibility from equation (2), where the outcome is the average share of his sons who died in specific ZIP rank quintile according to mean AGI (x-axis), either in the father's 1930 CZ of residence or outside it (panel). All regressions include fixed effects for father's occupation $\times$ year of birth, Social Security coverage category, race, and state/country of birth. Observations weighted using IPWs. 95% confidence intervals based on standard errors clustered at the level of father's industry of employment.

Figure A.15: Model performance: Precision-recall curves



Notes: This figure plots the performance of each model (panel title), as measured by the cross-validated precision-recall curve in the training set. Each dot corresponds to the precision and recall rate associated with a specific probability threshold.

Table A.1: Representativeness of Numident death-census linked sample

	Mean			Difference (two-sided *t*-test)	
	Population (20%)	Linked sample (unweighted)	Linked sample (weighted)	Linked sample (unweighted)	Linked sample (weighted)
	(1)	(2)	(3)	(4)	(5)
Panel (a): Father characteristics in 1930					
Black	0.082	0.036	0.081	-0.046	-0.001
				[0]	[0.091]
Born in 1875–1879	0.297	0.264	0.298	-0.033	0.001
				[0]	[0.475]
Born in 1880–1884	0.377	0.385	0.377	0.008	0
				[0]	[0.766]
Born in 1885–1888	0.326	0.35	0.325	0.025	0
				[0]	[0.697]
Born abroad	0.311	0.306	0.312	-0.005	0.001
				[0]	[0.443]
Born in Northeast	0.193	0.204	0.193	0.011	0
				[0]	[0.943]
Born in Midwest	0.244	0.264	0.244	0.02	0
				[0]	[0.939]
Born in South	0.228	0.203	0.227	-0.025	-0.001
				[0]	[0.403]
Born in West	0.024	0.022	0.023	-0.001	0
				[0]	[0.662]
Literate	0.934	0.948	0.935	0.014	0.001
				[0]	[0.011]
Covered by Social Security in 1935	0.875	0.88	0.874	0.005	-0.001
				[0]	[0.035]
Covered by Social Security in 1950	0.125	0.12	0.126	-0.005	0.001
				[0]	[0.035]
White-collar worker	0.225	0.221	0.21	-0.004	-0.015
				[0]	[0]
Farmer	0.004	0.004	0.004	0	0
				[0]	[0.251]
Skilled/semi-skilled worker	0.472	0.507	0.492	0.035	0.02
				[0]	[0]
Unskilled worker	0.299	0.268	0.294	-0.031	-0.006
				[0]	[0]
(table continues on next page)					
N	602,163		650,848		

Notes: Columns (3) and (5) weighted using IPWs. *p*-values for two-sided *t*-test of equality of means in brackets.

Table A.1 (cont.): Representativeness of Numident death-census linked sample

	Mean			Difference (two-sided *t*-test)	
	Population (20%)	Linked sample (unweighted)	Linked sample (weighted)	Linked sample (unweighted)	Linked sample (weighted)
	(1)	(2)	(3)	(4)	(5)
Panel (b): Household characteristics in 1930					
Live in Northeast	0.36	0.374	0.371	0.013	0.011
				[0]	[0]
Live in Midwest	0.312	0.336	0.32	0.024	0.008
				[0]	[0]
Live in South	0.214	0.197	0.214	-0.017	0.001
				[0]	[0.376]
Live in West	0.114	0.094	0.094	-0.02	-0.019
				[0]	[0]
Live in urban area	0.706	0.675	0.705	-0.031	-0.001
				[0]	[0.226]
Live on farm	0.065	0.076	0.065	0.012	0
				[0]	[0.752]
Homeowner	0.478	0.538	0.481	0.06	0.003
				[0]	[0.001]
Own a radio	0.438	0.469	0.44	0.03	0.002
				[0]	[0.088]
Number of sons	1.154	2.029	2.043	0.875	0.889
				[0]	[0]
Number of daughters	1.07	1.887	1.912	0.817	0.841
				[0]	[0]
N	602,163		650,848		

Notes: Columns (3) and (5) weighted using IPWs. *p*-values for two-sided *t*-test of equality of means in brackets.

Table A.2: Representativeness of Numident claims-census linked sample

	Mean			Difference (two-sided *t*-test)	
	Population (20%)	Linked sample (unweighted)	Linked sample (weighted)	Linked sample (unweighted)	Linked sample (weighted)
	(1)	(2)	(3)	(4)	(5)
Panel (a): Own characteristics in 1930					
Black	0.086	0.03	0.084	-0.056	-0.001
				[0]	[0.055]
Born in 1875–1879	0.297	0.338	0.29	0.041	-0.007
				[0]	[0]
Born in 1880–1884	0.377	0.392	0.375	0.015	-0.002
				[0]	[0.047]
Born in 1885–1888	0.326	0.27	0.336	-0.056	0.009
				[0]	[0]
Born abroad	0.311	0.267	0.311	-0.045	-0.001
				[0]	[0.453]
Born in Northeast	0.196	0.236	0.196	0.04	0.001
				[0]	[0.291]
Born in Midwest	0.245	0.305	0.245	0.061	0
				[0]	[0.622]
Born in South	0.224	0.164	0.223	-0.059	-0.001
				[0]	[0.546]
Born in West	0.025	0.028	0.025	0.003	0
				[0]	[0.831]
Literate	0.933	0.966	0.934	0.033	0.001
				[0]	[0.111]
Covered by Social Security in 1935	0.852	0.92	0.894	0.068	0.041
				[0]	[0]
Covered by Social Security in 1950	0.148	0.08	0.106	-0.068	-0.041
				[0]	[0]
White-collar worker	0.213	0.244	0.207	0.032	-0.005
				[0]	[0]
Farmer	0.003	0.002	0.003	-0.001	0
				[0]	[0.812]
Skilled/semi-skilled worker	0.448	0.521	0.497	0.074	0.049
				[0]	[0]
Unskilled worker	0.336	0.232	0.293	-0.104	-0.044
				[0]	[0]
(table continues on next page)					
N	758,021		364,278		

Notes: Columns (3) and (5) weighted using IPWs. *p*-values for two-sided *t*-test of equality of means in brackets.

Table A.2 (cont.): Representativeness of Numident claims-census linked sample

	Mean			Difference (two-sided *t*-test)	
	Population (20%) (1)	Linked sample (unweighted) (2)	Linked sample (weighted) (3)	Linked sample (unweighted) (4)	Linked sample (weighted) (5)
Panel (b): Household characteristics in 1930					
Live in Northeast	0.355	0.367	0.357	0.012	0.002
				[0]	[0.106]
Live in Midwest	0.316	0.356	0.323	0.04	0.006
				[0]	[0]
Live in South	0.208	0.151	0.196	-0.057	-0.011
				[0]	[0]
Live in West	0.121	0.126	0.124	0.005	0.003
				[0]	[0]
Live in urban area	0.701	0.741	0.699	0.04	-0.001
				[0]	[0.153]
Live on farm	0.082	0.053	0.084	-0.03	0.002
				[0]	[0.002]
Homeowner	0.466	0.539	0.466	0.073	0.001
				[0]	[0.444]
Own a radio	0.428	0.527	0.428	0.1	0
				[0]	[0.943]
N	758,021	364,278			

Notes: Columns (3) and (5) weighted using IPWs. *p*-values for two-sided *t*-test of equality of means in brackets.

Table A.3: Responses to 1940 Social Security Number question by Social Security coverage year

	Share of male workers in 1940 (%):	
	Without an SSN	With an SSN
	(1)	(2)
Panel (a): All employment categories, male workers aged 25–64		
Covered in 1935 based on 1940 employment	11.6	88.4
Covered in 1950 or later based on 1940 employment	71.4	28.6
Panel (b): Subset of employment categories used in the analysis, male workers aged 51–64		
Covered in 1935 based on 1940 employment	15.8	84.2
Covered in 1950 based on 1940 employment	69.6	30.4

Notes: Table shows the share of male workers reporting having an SSN vs. not in the 1940 Census (sample-line respondents only), separately for workers whose employment was covered by Social Security in 1935 vs. 1950 or later. Sample restricted to male workers aged 25–64 in 1940 in panel (a). Panel (b) further restricts the sample to the population of interest in the analysis: male workers aged 51–64 whose employment was covered in 1935 or 1950, excluding the self-employed.

Table A.4: Top 20 occupations with some Social Security coverage year variation in 1930

	Summary statistics among men born in 1875–1888 in 1930 Census						
	Population			Death-census linked sample		Claims-census linked sample	
Occupation	Number (1)	Share covered in 1935 (%) (2)	Share (%) (3)	Share covered in 1935 (%) (4)	Share (%) (5)	Share covered in 1935 (%) (6)	Share (%) (7)
Laborers (n.e.c.)	461,594	95.5	15.33	95.5	15.49	97	15.06
Janitors and sextons	49,589	54.9	1.65	48.3	1.68	66.6	1.5
Truck and tractor drivers	47,572	95.9	1.58	96	1.7	96.9	1.58
Clerical and kindred workers (n.e.c.)	44,365	93.1	1.47	93	1.3	95.7	1.56
Stationary engineers	33,251	93.3	1.1	93.3	1.15	95.3	1.14
Plumbers and pipe fitters	29,795	98.4	0.99	98.3	0.99	99	0.99
Electricians	25,574	96.2	0.85	95.9	0.83	97.9	0.9
Cooks, except private household	21,298	97.1	0.71	96.7	0.42	97.3	0.78
Taxicab drivers and chauffers	18,956	35.8	0.63	35.8	0.58	35	0.6
Guards, watchmen, and doorkeepers	17,628	88.6	0.59	89.2	0.6	92.2	0.55
Mechanics and repairmen (n.e.c.)	16,492	98.3	0.55	98.2	0.56	99	0.58
Stationary firemen	16,286	91.6	0.54	92.1	0.61	94.2	0.56
Accountants and auditors	16,082	94.3	0.53	94.4	0.47	96.4	0.52
Bookkeepers	15,983	97.3	0.53	97.5	0.5	97.4	0.57
Gardeners, except farm and groundskeepers	16,030	44.3	0.53	45.3	0.53	48.4	0.45
Waiters and waitresses	14,003	99	0.47	99.5	0.34	99.4	0.59
Porters	12,674	96	0.42	96.5	0.37	96.5	0.47
Teamsters	10,940	93.5	0.36	93.5	0.41	94.6	0.32
Engineers, civil	9,861	79	0.33	77	0.32	83.6	0.24
Laundry and dry cleaning operatives	7,322	98.5	0.24	98.3	0.21	99	0.25

Notes: Table displays summary statistics for the top 20 occupations with some Social Security coverage year variation in the 1930 Census (share covered in 1935 between 1 percent and 99 percent). Sample restricted to men born in 1875–1888 and whose current employment was covered in 1935 or 1950 (excluding the self-employed). Columns 4 and 5 restrict to the subset of men in our Numident death-census linked sample. Columns 6 and 7 restrict to the subset of men of our Numident claims-census linked sample. Individuals in columns 4–7 weighted using IPWs.

Table A.5: Selection of 1930 households and children into our linked sample, by children cohorts

	Dependent variable: Any/number of children linked							
	Children born in 1895–1930		Children born in 1895–1910		Children born in 1911–1920		Children born in 1921–1930	
	Any	Number	Any	Number	Any	Number	Any	Number
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel (a): All children								
Min. age of Social Security	-0.0006	-0.0014	0.0005	0.0006	-0.0008	-0.0016*	-0.0004	-0.0005
eligibility	(0.0006)	(0.0013)	(0.0005)	(0.0006)	(0.0006)	(0.0009)	(0.0007)	(0.0009)
Mean of dep. var.	0.216	0.362	0.036	0.039	0.174	0.241	0.065	0.083
R^2	0.0188	0.0178	0.0132	0.0131	0.0169	0.0157	0.0160	0.0158
Panel (b): Sons								
Min. age of Social Security	-0.0006	-0.0012	0.0001	0.0001	-0.0004	-0.0008	-0.0004	-0.0005
eligibility	(0.0006)	(0.0009)	(0.0002)	(0.0002)	(0.0005)	(0.0007)	(0.0004)	(0.0004)
Mean of dep. var.	0.136	0.178	0.015	0.016	0.1	0.118	0.039	0.045
R^2	0.0126	0.0118	0.0068	0.0068	0.0101	0.0096	0.0102	0.0099
Panel (c): Daughters								
Min. age of Social Security	-0.0002	-0.0002	0.0005	0.0006	-0.0007**	-0.0008**	-0.0000	-0.0000
eligibility	(0.0003)	(0.0005)	(0.0004)	(0.0004)	(0.0003)	(0.0004)	(0.0005)	(0.0005)
Mean of dep. var.	0.137	0.184	0.021	0.023	0.103	0.123	0.034	0.038
R^2	0.0126	0.0124	0.0084	0.0083	0.0106	0.0102	0.0095	0.0093
N	2,971,208	2,971,208	2,971,208	2,971,208	2,971,208	2,971,208	2,971,208	2,971,208

Notes: Unit of observation is a 1930 household headed by a man born in 1875–1888 and covered by Social Security in 1935 or 1950 based on his employment information, excluding self-employed workers. The dependent variable is either an indicator equal to one if at least one child of a specific sex and cohort range was linked to the 1930 household, or a discrete variable capturing the number of children of a specific sex and cohort range linked to the 1930 household. All regressions include fixed effects for head of household occupation × year of birth, Social Security coverage category, race, and state/country of birth. Robust standard errors in parentheses, clustered at the head of household industry level. *** 1%, ** 5%, * 10% significance.

Table A.6: Numident death-census linked sample: summary statistics

	Mean outcomes		
	All children (1895–1930) (1)	Sons (1895–1930) (2)	Daughters (1895–1930) (3)
Panel (a): Migration			
Died outside parents' 1930 county of residence	0.626	0.629	0.623
Died outside parents' 1930 county of residence, but within parents' 1930 commuting zone of residence	0.143	0.139	0.147
Died outside parents' 1930 commuting zone of residence	0.482	0.489	0.476
Died outside parents' 1930 state of residence	0.361	0.37	0.353
Died outside parents' 1930 census division of residence	0.275	0.287	0.265
Died outside parents' 1930 census region of residence	0.23	0.243	0.219
Panel (b): Death ZIP characteristics			
Percentile rank based on mean 2001 adjusted gross income (AGI)	62.6	62.7	62.4
Percentile rank based on median 2000 house values	65.9	66.2	65.7
Mean 2001 AGI (2001 dollars)	46,487	46,552	46,425
Median 2000 house values (2000 dollars)	138,672	139,123	138,242
Number of children	1,091,106	536,670	554,436
Number of families	650,848	409,326	411,524

Notes: Table displays summary statistics in our Numident death-census linked sample.

Table A.7: OCG sample: summary statistics

	Mean outcomes	
	Sons (1907–1930) (1)	Sons (1907–1919) (2)
Panel (a): Income (1972 dollars, unless otherwise specified)		
Social Security earnings, 1937–1950 (nominal)	13,504	17,095
Social Security earnings, 1951–1960	45,980	46,719
Social Security earnings, 1961–1970	58,548	57,274
Social Security earnings, 1971–1975	30,334	27,492
Adjusted gross income (AGI), 1972	14,813	14,157
IRS wage and salary, 1972	12,341	11,298
Panel (b): Occupational standing (1950 dollars)		
Occupational income score (1950 basis), first job	2,561	2,510
Occupational income score (1950 basis), 1962 job	3,074	3,076
Occupational income score (1950 basis), 1973 job	3,162	3,132
N	3,020	1,739

Notes: Table displays summary statistics in our OCG sample.

Table A.8: Numident death-census linking: Stage 1 and stage 2 features

Stage 1 set-level features	
Type	Description
Continuous	Gap between highest and second-highest JW score
Continuous	Max JW score
Continuous	Mean JW score
Discrete	Absolute difference in number of matching kids (method 1) between highest and second-highest JW score potential
Continuous	Gap between highest and second-highest dad last JW score
Continuous	Max dad last JW score
Continuous	Mean dad last JW score
Discrete	Absolute difference in number of matching kids (method 1) between highest and second-highest dad last JW score potential
Binary	Unique potential with 1+ matching kids (method 1)
Binary	Unique potential with 2+ matching kids (method 1)
Binary	Unique potential with max JW score or max dad last JW score & 1+ matching kids (method 1)
Binary	Unique potential with max JW score or max dad last JW score & 2+ matching kids (method 1)
Binary	Unique potential with max JW score or max dad last JW score & dad or mom middle initial match & 1+ matching kids (method 1)
Binary	Unique potential with max JW score or max dad last JW score & dad or mom middle initial match & 2+ matching kids (method 1)
Binary	Unique potential with 1+ matching kids (method 1) & 1 matching kid (method 2)
Binary	Unique potential with 1+ matching kids (method 1) & 2+ matching kids (method 2)
Binary	Unique potential with max JW score or max dad last JW score & 1+ matching kids (method 1) & 1 matching kid (method 2)
Binary	Unique potential with max JW score or max dad last JW score & 1+ matching kids (method 1) & 2+ matching kids (method 2)
Binary	Unique potential with max JW score or max dad last JW score & dad or mom middle initial match & 1+ matching kids (method 1) & 1 matching kid (method 2)
Binary	Unique potential with max JW score or max dad last JW score & dad or mom middle initial match & 1+ matching kids (method 1) & 2+ matching kids (method 2)
Binary	Unique potential with max JW score or max dad last JW score & 1+ matching kids (method 1) or 1+ matching kids (method 2) & dad or mom middle initial match & 1+ matching kids (method 3)
Binary	Unique potential with max JW score or max dad last JW score & 1+ matching kids (method 1) or 1+ matching kids (method 2) & dad or mom middle initial match & 1+ matching kids (method 3)
Binary	Unique potential with max JW score or max dad last JW score & 1+ matching kids (method 1) or 1+ matching kids (method 2) & dad or mom middle initial match & 1+ matching kids (method 4)
Binary	Unique potential with max JW score or max dad last JW score & 1+ matching kids (method 1) or 1+ matching kids (method 2) & dad or mom middle initial match & 1+ matching kids (method 4)

Stage 2 pair-level features			
Type	Description	Type	Description
Continuous	JW score	Binary	Modal race match
Continuous	Dad last JW score	Binary	Dad middle initial match
Discrete	Number of matching kids (method 1)	Binary	Mom middle initial match
Discrete	Number of matching kids (method 2)	Binary	Max JW score or max dad last JW score & 1+ matching kids (method 1)
Discrete	Number of matching kids (method 3)	Binary	Max JW score or max dad last JW score & 1+ matching kids (method 2)
Discrete	Number of matching kids (method 4)	Binary	Max JW score or max dad last JW score & 1+ matching kids (method 3)
Discrete	Rank	Binary	Max JW score or max dad last JW score & 1+ matching kids (method 4)
Binary	Rank = 1		

Notes: "JW score" is the average of three Jaro-Winkler scores comparing primary and potential: (1) father first names, (2) father last names, and (3) mother first names. Two kids match according to method 1 if three conditions hold: (1) the Jaro-Winkler score between their first names is above 0.8 or if the Levenshtein distance between their first names is within 2, (2) their absolute age difference is within 2 years, and (3) their birthplaces match. Method 2 allows for inverted first and middle names, method 3 allows for matches on initials (both first and middle name initials must match), and method 4 is a combination of methods 2 and 3.